



US009313526B2

(12) **United States Patent**
Bivolarsky et al.

(10) **Patent No.:** **US 9,313,526 B2**
(45) **Date of Patent:** **Apr. 12, 2016**

(54) **DATA COMPRESSION FOR VIDEO**

(75) Inventors: **Lazar Bivolarsky**, Cupertino, CA (US);
Renat Vafin, Tallinn (EE); **Mattias Nilsson**, Sundbyberg (SE); **Soren Vang Andersen**, Luxembourg (LU)

(73) Assignee: **Skype**, Dublin (IE)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 556 days.

(21) Appl. No.: **13/029,710**

(22) Filed: **Feb. 17, 2011**

(65) **Prior Publication Data**

US 2012/0044990 A1 Feb. 23, 2012

Related U.S. Application Data

(60) Provisional application No. 61/306,385, filed on Feb. 19, 2010, provisional application No. 61/365,242, filed on Jul. 16, 2010.

(51) **Int. Cl.**

H04N 19/00 (2014.01)
H04N 19/48 (2014.01)
H04N 19/146 (2014.01)
H04N 19/85 (2014.01)
H04N 19/105 (2014.01)
H04N 19/176 (2014.01)
H04N 19/46 (2014.01)

(Continued)

(52) **U.S. Cl.**

CPC **H04N 19/85** (2014.11); **H04N 19/105** (2014.11); **H04N 19/126** (2014.11); **H04N 19/132** (2014.11); **H04N 19/146** (2014.11); **H04N 19/154** (2014.11); **H04N 19/176** (2014.11); **H04N 19/46** (2014.11); **H04N 19/48** (2014.11); **H04N 19/60** (2014.11); **H04N 19/61** (2014.11)

(58) **Field of Classification Search**

USPC 375/240.03, E7.2, E7.214,
375/240.02–240.29, E7.176, E7.189;
714/776; 382/250–260

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

1,691,147 A * 11/1928 Clark et al. 358/400
4,833,535 A 5/1989 Ozeki et al.

(Continued)

FOREIGN PATENT DOCUMENTS

CN 1109245 9/1995
CN 1256049 6/2000

(Continued)

OTHER PUBLICATIONS

Bennett, "Spectra of Quantized Signals," Bell System Technical Journal, vol. 27, Jul. 1948, pp. 446-472.*

(Continued)

Primary Examiner — Jay Patel



Assistant Examiner — Reza Aghevli

(74) *Attorney, Agent, or Firm* — Danielle Johnston-Holmes; Micky Minhas

(57) **ABSTRACT**

The present invention provides a technique for performing one or more aspects of video coding such as quantization, intra prediction coding or inter prediction coding in dependence on a perceptual model taking into account human sensitivity to data in the video signal. The perceptual model may relate to spatial frequency, temporal frequency, contrast sensitivity, color sensitivity, a structural metric, and/or one or more parameters affecting perception such as motion in the video, the distance of a recipient user from the screen, and the size, aspect ratio or resolution of the screen of the recipient terminal.

46 Claims, 27 Drawing Sheets



(51)	Int. Cl.				
	H04N 19/61	(2014.01)	2011/0206110	A1	8/2011 Bivolarsky et al.
	H04N 19/60	(2014.01)	2011/0206113	A1	8/2011 Bivolarsky et al.
	H04N 19/126	(2014.01)	2011/0206117	A1	8/2011 Bivolarsky et al.
	H04N 19/132	(2014.01)	2011/0206118	A1	8/2011 Bivolarsky et al.
	H04N 19/154	(2014.01)	2011/0206119	A1	8/2011 Bivolarsky et al.
			2011/0206131	A1	8/2011 Vafin et al.
			2011/0206132	A1	8/2011 Bivolarsky et al.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,719,642	A	2/1998	Lee	
5,732,157	A	3/1998	Osawa	
5,790,206	A	8/1998	Ju	
6,091,767	A	7/2000	Westerman	
6,091,777	A *	7/2000	Guetz et al.	375/240.11
6,148,109	A *	11/2000	Boon et al.	382/238
6,201,834	B1	3/2001	Zhu	
6,408,026	B1 *	6/2002	Tao	375/240.03
6,445,810	B2 *	9/2002	Darrell et al.	382/115
6,650,782	B1 *	11/2003	Joshi et al.	382/239
6,668,020	B2	12/2003	Ma et al.	
6,856,651	B2	2/2005	Singh	
7,430,329	B1	9/2008	Sarna	
8,681,873	B2	3/2014	Bivolarsky et al.	
8,913,661	B2	12/2014	Bivolarsky et al.	
9,078,009	B2	7/2015	Bivolarsky et al.	
2001/0002205	A1	5/2001	Beattie	
2001/0014124	A1	8/2001	Nishikawa	
2002/0012396	A1	1/2002	Pau et al.	
2003/0001757	A1	1/2003	Kondo et al.	
2003/0016751	A1 *	1/2003	Vetro et al.	375/240.16
2003/0026340	A1	2/2003	Divakaran et al.	
2003/0067981	A1 *	4/2003	Zhao et al.	375/240.08
2003/0081678	A1	5/2003	Fukuda	
2003/0179825	A1	9/2003	Sekiguchi et al.	
2003/0206588	A1	11/2003	Etoh et al.	
2004/0021592	A1	2/2004	Karczewicz	
2004/0057516	A1 *	3/2004	Kim et al.	375/240.16
2004/0170395	A1 *	9/2004	Filippini et al.	386/111
2004/0213348	A1	10/2004	Kim et al.	
2004/0228540	A1	11/2004	Chen et al.	
2005/0041740	A1 *	2/2005	Sekiguchi et al.	375/240.16
2005/0094727	A1	5/2005	Ha et al.	
2005/0229200	A1 *	10/2005	Kirkland et al.	725/12
2006/0013303	A1	1/2006	Nguyen et al.	
2006/0104352	A1	5/2006	Chen et al.	
2006/0120452	A1	6/2006	Li	
2006/0152597	A1	7/2006	Shimizu	
2006/0153295	A1	7/2006	Wang et al.	
2006/0239354	A1	10/2006	Amano et al.	
2007/0104381	A1	5/2007	Chen et al.	
2007/0116370	A1	5/2007	Smirnov	
2007/0120967	A1	5/2007	Eshkoli et al.	
2007/0172135	A1	7/2007	Song	
2007/0177671	A1	8/2007	Yang	
2007/0217511	A1	9/2007	Li et al.	
2007/0253480	A1 *	11/2007	Tsujii et al.	375/240.08
2008/0037642	A1	2/2008	Tsuchiya	
2008/0075171	A1	3/2008	Suzuki	
2008/0095230	A1 *	4/2008	Hannuksela et al.	375/240.02
2008/0123750	A1 *	5/2008	Bronstein et al.	375/240.24
2008/0192838	A1 *	8/2008	Chen et al.	375/240.25
2008/0240247	A1	10/2008	Lee et al.	
2008/0273810	A1	11/2008	Subramania et al.	
2009/0003437	A1 *	1/2009	Cho et al.	375/240.03
2009/0060362	A1	3/2009	Harmanci et al.	
2009/0067503	A1	3/2009	Jeong et al.	
2009/0080788	A1	3/2009	Lynch et al.	
2009/0080804	A1	3/2009	Hamada et al.	
2009/0097548	A1	4/2009	Karczewicz et al.	
2009/0110054	A1 *	4/2009	Kim et al.	375/240.1
2009/0180538	A1	7/2009	Visharam et al.	
2009/0225830	A1	9/2009	He et al.	
2009/0257669	A1	10/2009	Kim et al.	
2010/0080297	A1	4/2010	Wang et al.	
2010/0166073	A1	7/2010	Schmit et al.	
2010/0310065	A1	12/2010	Chang et al.	

FOREIGN PATENT DOCUMENTS

CN	1262496	8/2000
CN	1266589	9/2000
CN	1805544	7/2006
CN	101061722	10/2007
CN	101133650	2/2008
CN	101511014	8/2009
CN	101578769	11/2009
CN	101584120	11/2009
EP	0 639 925	2/1995
EP	0639925	2/1995
EP	0 679 033	10/1995
EP	0679033	10/1995
EP	1024456	8/2000
EP	1024456	8/2000
JP	2007-60437	3/2007
JP	200760437	3/2007
JP	2009509408	3/2009
WO	WO 00/27128	5/2000
WO	WO-0027128	5/2000
WO	WO 2007/011851	1/2007
WO	WO-2007011851	1/2007

OTHER PUBLICATIONS

Smith, "Instantaneous Companding of Quantized Signals," Bell System Technical Journal, vol. 36, May 1957, pp. 653-709.*

Kaneko, "A Unified Formulation of Segment Companding Laws and Synthesis of Codecs and Digital Companders," Bell System Technical Journal, vol. 49, Sep. 1970, pp. 1555-1558.*

Brokish et al., "A-Law and mu-Law Companding Implementation Using the TMS320C54x," Texas Instruments, Application Note: SPR163A, Dec. 1997.*

Sullivan et al., "Rate-Distortion Optimization for Video Compression", IEEE Signal Processing Magazine, Nov. 1998, pp. 74-90.*

Wiegand et al., "Overview of the H.264/AVC Video Coding Standard", IEEE Trans. Cir. Sys. Video Technology, vol. 13, No. 7, Jul. 2003, pp. 560-576.*

Sullivan et al., "Video Compression—From Concepts to the H.264/AVC Standard", Proc. IEEE, vol. 93, No. 1, Jan. 2005, pp. 18-31.*

Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority, or the Declaration, PCT/EP2011/052458, Mailing Date: Oct. 14, 2011, 32 pp.

Cavallaro, A. et al., "Perceptual prefiltering for video coding," Intelligent Multimedia, Video and Speech Processing, 2004. Proceedings of 2004 International Symposium on Hong Kong, China, Oct. 20-22, 2004. Piscataway, NJ, USA IEEE, pp. 510-513, XP010801531, Oct. 20, 2004.

Song, M. et al., "Motion estimation in DCT domain", Communication Technology Proceedings, 1996. ICCT '96. 1996 International AI Conference on Beijing, China May 5-7, 1996, New York, NY, USA, IEEE, vol. 2, pp. 670-674, XP010532901, May 5, 1996.

Lin, Y.C. et al., "Fast Full-Search Block-Matching Algorithm for Motion-Compensated Video Compression", IEEE Transactions on Communications, IEEE Service Center, Piscataway, NJ, USA, vol. 45, No. 5, XP011008939, May 1, 1997.

"Non-Final Office Action", U.S. Appl. No. 12/838,197, (Dec. 7, 2012), 14 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,139, (Dec. 7, 2012), 13 pages.

"Final Office Action", U.S. Appl. No. 12/838,131, (Jan. 30, 2013), 15 pages.

"Final Office Action", U.S. Appl. No. 12/838,124, (Feb. 15, 2013), 19 pages.

(56)

References Cited**OTHER PUBLICATIONS**

- "Preliminary Examination Report", Application No. PCT/EP2011/052465, (Jun. 5, 2012), 10 pages.
- "Information Technology—JPEG 2000 image coding system: Core coding system", IAO 15444 ITU-T Recommendation T.800, (Aug. 2002), 212 pages.
- International Search Report and Written Opinion, PCT/EP2011/052462, Date of Mailing: May 24, 2011.
- Richardson, I.E., "H.264 and MPEG-4," H.264 and MPEG-4 Video Compression: Video Coding for Next-generation Multimedia, Chapter 6, Part 10, XP030001626, ISBN: 0-470-84837-5 (Oct. 17, 2003).
- Freina, T., et al., "Predictive Fractal Image Coding: Hybrid Algorithms and Compression of Residuals," 1998 Data Compression Conference, p. 549 (Mar. 30, 1998).
- Vasconcelos, N., et al., "Library-based image coding," Proceedings of the International Conference on Acoustics, Speech and Signal Processing, 19: 489-492 (Apr. 19, 1994).
- Vasconcelos, N., et al., "Library-based coding: a representation for efficient video compression and retrieval," Data Compression Conference, pp. 121-130 (Mar. 25, 1997).
- Invitation to Pay Additional Fees and Partial International Search Report, PCT/EP2011/052464, Date of Mailing: Jun. 29, 2011.
- Schultz, R., et al., "Subpixel Motion Estimation for Super-Resolution Image Sequence Enhancement," Journal of Visual Communication and Image Representation, vol. 9(1), pp. 38-50 (1998).
- Ben-Ezra, M., et al., "Video Super-Resolution Using Controlled Subpixel Detector Shifts," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 27(6), pp. 977-987 (2005).
- Segall, A., et al., "Improved High-Definition Video by Encoding at an Intermediate Resolution," Visual Communications and Image Processing, SPIE vol. 5308, pp. 1007-1018 (2004).
- Barreto, D., et al., "Region-based Super-Resolution for Compression," Multidim Syst Sign Process, vol. 18, pp. 59-81 (2007).
- Molina, R., et al., "Toward a New Video Compression Scheme Using Super-Resolution," Visual Communications and Image Processing, SPIE vol. 6077 (2006).
- Invitation to Pay Additional Fees and, Where Applicable, Protest Fee, dated Jul. 5, 2011, for Application No. PCT/EP2011/052465, consisting of 11 pages.
- International Search Report and Written Opinion, PCT/EP2011/052464, date of mailing Sep. 1, 2011.
- Hoang, D.T., et al., "Efficient cost measures for motion compensation at low bit rates," Data Compression Conference, pp. 102-111 (Mar. 31, 1996).
- Sullivan, G.J., et al., "Rate-Distortion Optimization for Video Compression," IEEE Signal Processing Magazine, 15(6): 74-90 (Nov. 1, 1998).
- Hayes, M.H., "Iterated Function Systems for Image and Video Coding," Journal on Communications, 45: 11-19 (May 1, 1994).
- Rodrigues, N.M.M., et al., "Improving multiscale recurrent pattern image coding with enhanced dictionary updating strategies," Telecommunications Symposium, pp. 257-262 (Sep. 1, 2006).
- Chang, S., et al., "Manipulation and compositing of MC-DCT compressed video," Journal on Selected Areas in Communications, 13(1): 1-11 (Jan. 1, 1995).
- Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority for Int'l Application No. PCT/EP2011/052465; Date of Mailing: Sep. 26, 2011.
- Vitter, J. S., "Design and Analysis of Dynamic Huffman Codes," Journal of the Association for Computing Machinery, vol. 34, No. 4, pp. 825-845 (1987).
- Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority, dated Apr. 20, 2011, for Application No. PCT/EP2011/052449, consisting of 9 pages.
- Mingzhou, S., et al., "Motion Estimation in DCT Domain," IEEE, vol. 5, pp. 670-674 (1996).
- Richardson, I. E., "H.264 and MPEG-4 Video Compression. Chapter 7: Design and Performance," pp. 225-267 (2003).
- Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority, or the Declaration, for Application No. PCT/EP2011/052460, mailed on May 24, 2011, consisting of 19 pages.
- Notification of Transmittal of the International Search Report and the Written Opinion of the International Searching Authority, or the Declaration, International Application: PCT/EP2011/052446, Mailing Date: Dec. 29, 2011, 13 pp.
- Ginzburg, A., et al., "DCT-Domain Coder for Digital Video Applications", 2006 International Conference on Information Technology: Research and Education, pp. 110-114 (Oct. 1, 2006).
- "Final Office Action", U.S. Appl. No. 12/838,118, (Mar. 15, 2013), 16 pages.
- "Final Office Action", U.S. Appl. No. 12/838,139, (Mar. 27, 2013), 12 pages.
- "Final Office Action", U.S. Appl. No. 12/838,168, (Apr. 12, 2013), 19 pages.
- "Final Office Action", U.S. Appl. No. 12/838,197, (May 9, 2013), 12 pages.
- "Notice of Allowance", U.S. Appl. No. 12/838,159, (Mar. 29, 2013), 11 pages.
- Lee, J-B and Kalva, H., "Intra Prediction," in *The VC-1 and H.264 Video Compression Standards for Broadband Video Services* (NY: Springer Science+Business Media), pp. 247-277 (2010).
- Weigand, T., et al., Overview of the H.264/AVC Video Coding Standard, *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7): 560-576 (Jul. 2003).
- Bivolarski, L.M., et al., "Design Considerations for Computationally Constrained Two-Way Real-Time Video Communication," *Applications of Digital Image Processing XXXII*, A.G. Tescher, ed., in *Proc. of SPIE*, 7443: 74430I-1-74430I-17 (2009).
- Xiao, Feng, "DCT-based Video Quality Evaluation," Final Project for EE392J, Winter 2000.
- "Final Office Action", U.S. Appl. No. 12/838,159, (Jul. 24, 2012), 15 pages.
- "Non-Final Office Action", U.S. Appl. No. 12/838,131, (Aug. 30, 2012), 13 pages.
- "Non-Final Office Action", U.S. Appl. No. 12/838,168, (Aug. 31, 2012), 16 pages.
- "Non-Final Office Action", U.S. Appl. No. 12/838,124, (Sep. 6, 2012), 18 pages.
- "Non-Final Office Action", U.S. Appl. No. 12/838,118, (Sep. 14, 2012), 16 pages.
- "Restriction Requirement", U.S. Appl. No. 12/838,139, (Oct. 1, 2012), 6 pages.
- "Restriction Requirement", U.S. Appl. No. 12/838,197, (Oct. 16, 2012), 6 pages.
- Invitation to Pay Additional Fees and, Where Applicable, Protest Fee and Partial International Search for Int'l Application No. PCT/EP2011/052458; Date of Mailing: May 31, 2011, 9 pp.
- Hong, G. et al., "SenseCoding: Accelerometer-Assisted Motion Estimation for Efficient Video Encoding," Proceedings of the 2008 ACM International Conference on Multimedia with Co-Located Symposium & Workshops: Area'08, Communicability'08, HCC'08, MIR'08, MS'08, SAME'08, SRMC'08, TVS'08, VNBA'08: Vancouver, BC, Canada, Oct. 27-31, 2008, ACM Press, Oct. 26, 2008, pp. 749-752, XP007918610.
- Lee, S.-W. et al., "Video Rate Control Using an Adaptive Quantization Based on a Combined Activity Measure," Aug. 17, 2005, Knowledge-Based Intelligent Information and Engineering Systems; [Lecture Notes in Computer Science; Lecture Notes in Artificial Intelligence; LNCS], Springer-Verlag, Berlin/Heidelberg, pp. 1297-1302, XP019015301.
- Rosenholtz, R. et al., "Perceptual adaptive JPEG coding," Proceedings of the International Conference on Image Processing (ICIP) Lausanne, Sep. 16-19, 1996, New York, IEEE, US, vol. 1, Sep. 16, 1996, pp. 901-904, XP010202284.
- Tang, C.-W. et al., "Spatiotemporal Visual Considerations for Video Coding," IEEE Transactions on Multimedia, IEEE Service Center, Piscataway, NJ, US, vol. 9, No. 2, Feb. 1, 2007, pp. 231-238, XP011157485.

(56)

References Cited

OTHER PUBLICATIONS

Wallace, G.K., "Overview of the JPEG (ISO/CCITT) Still Image Compression Standard," Joint Photographic Expert Group Conference, Crowborough: JPEG Forum Ltd., GB, Nov. 1, 1989, XP017203197.

"Final Office Action", U.S. Appl. No. 12/838,131, (Nov. 26, 2013), 13 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,131, (Aug. 1, 2013), 15 pages.

"Notice of Allowance", U.S. Appl. No. 12/838,159, (Nov. 13, 2013), 12 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,131, Feb. 27, 2014, 18 pages.

"Extended European Search Report", EP Application No. 13166503.6, Nov. 27, 2013, 7 pages.

"Final Office Action", U.S. Appl. No. 12/838,124, Jun. 5, 2014, 19 pages.

"Final Office Action", U.S. Appl. No. 12/838,131, Jun. 6, 2014, 9 pages.

"Foreign Office Action", CN Application No. 201180010061.0, Aug. 11, 2014, 17 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,118, Aug. 1, 2014, 18 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,139, Jul. 31, 2014, 16 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,168, Aug. 1, 2014, 16 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,197, Aug. 14, 2014, 14 pages.

"Notice of Allowance", U.S. Appl. No. 12/838,131, Aug. 11, 2014, 7 pages.

"Final Office Action", U.S. Appl. No. 12/838,118, Jan. 22, 2015, 17 pages.

"Foreign Office Action", CN Application No. 201180010045.1, Oct. 8, 2014, 22 pages.

"Foreign Office Action", EP Application No. 11705195.3, Dec. 3, 2013, 5 pages.

"Notice of Allowance", U.S. Appl. No. 12/838,197, Mar. 2, 2015, 8 pages.

"Final Office Action", U.S. Appl. No. 12/838,124, Jun. 2, 2015, 23 pages.

"Final Office Action", U.S. Appl. No. 12/838,139, May 20, 2015, 17 pages.

"Foreign Office Action", CN Application No. 201180010045.1, Apr. 30, 2015, 12 pages.

"Foreign Office Action", CN Application No. 201180010048.5, May 15, 2015, 10 pages.

"Foreign Office Action", CN Application No. 201180010061.0, Jun. 12, 2015, 10 pages.

"Foreign Office Action", CN Application No. 201180010082.2, Sep. 2, 2015, 8 pages.

"Foreign Office Action", CN Application No. 201180010090.7, Sep. 9, 2015, 8 pages.

"Foreign Office Action", EP Application No. 11705195.3, Sep. 15, 2015, 4 pages.

"Supplemental Notice of Allowance", U.S. Appl. No. 12/838,197, Jun. 5, 2015, 2 pages.

"Final Office Action", U.S. Appl. No. 12/838,168, Jan. 2, 2015, 18 pages.

"Foreign Office Action", CN Application No. 201180010048.5, Nov. 2, 2014, 4 pages.

"Foreign Office Action", CN Application No. 201180010082.2, Nov. 25, 2014, 17 pages.

"Foreign Office Action", CN Application No. 201180010090.7, Nov. 3, 2014, 16 pages.

"Foreign Office Action", CN Application No. 201180010098.3, Dec. 2, 2014, 12 pages.

"Foreign Office Action", CN Application No. 201180010214.1, Nov. 25, 2014, 16 pages.

"Foreign Office Action", CN Application No. 201180010289.X, Dec. 16, 2014, 14 pages.

"Foreign Office Action", EP Application No. 11705195.3, Oct. 21, 2014, 5 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,124, Jan. 12, 2015, 20 pages.

Hoang, et al., "Efficient Cost Measures for Motion Estimation at Low Bit Rates", IEEE Transactions on Circuits and Systems for Video Technology vol. 8 No. 4, Aug. 31, 1998, 13 pages.

Song, "Research on Super Resolution Reconstruction of Video and Image Sequences", China Doctor Dissertation, Full-text Database, Collection of Information, Science and Technology—English Abstract pp. 8-9, Jan. 15, 2014, 44 pages.

Zhu, et al., "Transform-Exempted Calculation of Sum of Absolute Hadamard Transformed Differences", IEEE Transactions on Circuits and Systems for Video Technology archive vol. 19 Issue 8, Aug. 2009, 6 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052449, (Apr. 20, 2011), 9 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052460, (May 24, 2011), 19 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052446, (Dec. 29, 2011), 13 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052458, (Oct. 14, 2011), 32 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052464, (Sep. 1, 2011), 25 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052465, (Sep. 26, 2011), 21 pages.

"International Search Report and Written Opinion", International Application No. PCT/EP2011/052462, (May 24, 2011), 14 pages.

"Invitation to Pay Additional Fees and, Where Applicable, Protest Fee and Partial International Search", International Application No. PCT/EP2011/052458, (May 31, 2011), 9 pages.

"Invitation to Pay Additional Fees and, Where Applicable, Protest Fee and Partial International Search", International Application No. PCT/EP2011/052464, (Jun. 29, 2011), 11 pages.

"Invitation to Pay Additional Fees and, Where Applicable, Protest Fee and Partial International Search", International Application No. PCT/EP2011/052465, (Jul. 5, 2011), 10 pages.

"Non-Final Office Action", U.S. Appl. No. 12/838,159, (Apr. 13, 2012), 11 pages.

"PCT Search Report and Written Opinion", Application No. PCT/EP2011/052463, (Mar. 16, 2012), 13 pages.

Barreto, et al., "Region-Based Super-Resolution for Compression", *Multidim Syst Sign Process*, vol. 18, (2007), pp. 59-81.

Ben-Ezra, Moshe et al., "Video Super-Resolution Using Controlled Subpixel Detector Shifts", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27(6), (2005), pp. 977-987.

Bivolarski, Lazar M., et al., "Design Considerations for Computationally Constrained Two-Way Real-Time Video Communication", *Applications of Digital Image Processing XXXII*, A.G. Tescher, ed., in Proceedings of SPIE, 7443:7440I-1-74430II-17, (2009), 17 pages.

Cavallaro, Andrea et al., "Perceptual Prefiltering for Video Coding", *Intelligent Multimedia, Video and Speech Processing*, 2004, Proceedings of 2004 International Symposium on Hong Kong, China, Oct. 20-22, 2004. Piscataway, NJ, IEEE, XP010801531, (Oct. 20, 2004), pp. 510-513.

Chang, Shih-Fu et al., "Manipulation and Compositing of MC-DCT Compressed Video", *Journal on Selected Areas in Communications*, 13(1), (Jan. 1, 1995), pp. 1-11.

Freina, Thomas et al., "Predictive Fractal Image Coding: Hybrid Algorithms and Compression of Residuals", *1998 Data Compression Conference*, (Mar. 30, 1998), p. 549.

Ginzburg, Alex et al., "DCT-Domain Coder for Digital Video Applications", *International Conference on Information Technology: Research and Education*, (Oct. 1, 2006), pp. 110-114.

Hayes, "Iterated Function Systems for Image and Video Coding", *Journal on Communications*, 45, (May 1, 1994), pp. 11-19.

Hoang, Dzong T., et al., "Efficient Cost Measures for Motion Compensation at Low Bit Rates", *Data Compression Conference*, (Mar. 31, 1996), 10 pages.

(56)

References Cited

OTHER PUBLICATIONS

Hong, Guangming et al., "SenseCoding: Accelerometer-Assisted Motion Estimation for Efficient Video Encoding", *Proceedings of the 2008 ACM International Conference on Multimedia with Co-located Symposium & Workshops: Area '08, Communicabilityms '08, HCC '08, MIR '08, MS'08, SAME '08, SRMC '08, TVS '08, VNBA '08*: Vancouver, BC, Canada, Oct. 27-31, 2008 XP0079186, (Oct. 26, 2008), pp. 749-752.

Lee, Jae-Beom et al., "Intra Prediction", *The VC-1 and H.264 Video Compression Standards for Broadband Video Services*, NY: Springer Science+Business Media, LLC, (2010), pp. 247-277.

Lee, Si-Woong et al., "Video Rate Control Using an Adaptive Quantization Based on a Combined Activity Measure", *Knowledge-Based Intelligent Information and Engineering Systems; [Lecture Notes in Computer Science: Lecture Notes in Artificial Intelligence; LNCS]*, Springer-Verlag, Berlin/Heidelberg, XP019015301, (Aug. 17, 2005), pp. 1297-1302.

Lin, Yih-Chuan et al., "Fast Full-Search Block-Matching Algorithm for Motion-Compensated Video Compression", *IEEE Transactions on Communications*, IEEE Service Center, Piscataway, NJ, vol. 45, No. 5, XP011008939, (May 1, 1997), pp. 527-531.

Molina, et al., "Towards a New Video Compression Scheme Using Super-Resolution", *Visual Communications and Image Processing*, SPIE vol. 6077, (2006), 11 pages.

Nuno, "Improving Multiscale Recurrent Pattern Image Coding with Enhanced Dictionary Updating Strategies", *Telecommunications Symposium*, (Sep. 1, 2006), pp. 257-262.

Richardson, Iain E., "H.264 and MPEG-4 Video Compression. Chapter 7: Design and Performance", *John Wiley & Sons, Ltd.* ISBN: 0-470-84837-5, (2003), pp. 225-234.

Richardson, Iain E., "H.264/MPEG4 Part 10", *H.264 and MPEG-4 Video Compression: Video Coding for Next-Generation Multimedia*, John Wiley & Sons, Ltd., ISBN: 0-470-84837-5, (Oct. 17, 2003), 12 pages.

Rosenholtz, Ruth et al., "Perceptual Adaptive JPEG Coding" *Proceedings of the International Conference on Image Processing (ICIP)*, Lausanne, Sep. 16-19, 1996, New York, IEEE, US, vol. 1 XP010202284, (Sep. 16, 1996), pp. 901-904.

Schultz, Richard R., et al., "Subpixel Motion Estimation for Super-Resolution Image Sequence Enhancement", *Journal of Visual Communication and Image Representation*, vol. 9(1), (1998), pp. 38-50.

Segall, Andrew et al., "Improved High-Definition Video by Encoding at an Intermediate Resolution", *Visual Communications and Image Processing, SPIE* vol. 5308, (2004), pp. 1007-1018.

Song, Mingzhou et al., "Motion Estimation in DCT Domain", In *Proceedings of ICCT 1996*, (May 1996), pp. 670-674.

Sullivan, Gary J., et al., "Rate-Distortion Optimization for Video Compression", *IEEE Signal Processing Magazine*, 15(6), (Nov. 1, 1998), pp. 74-90.

Tang, Chih-Wei "Spatiotemporal Visual Considerations for Video Calling" *IEEE Transactions on Multimedia*, IEEE Service Center, Piscataway, NJ, US, vol. 9, No. 2, XP011157485, (Feb. 1, 2007), pp. 231-238.

Vasconcelos, Nuno et al., "Library-Based Coding: A Representation for Efficient Video Compression and Retrieval", *Data Compression Conference*, (Mar. 25, 1997), 10 pages.

Vasconcelos, Nuno et al., "Library-Based Image Coding", *Proceedings of the International Conference on Acoustics, Speech and Signal Processing*, 19, (Apr. 19, 1994), pp. 489-492.

Vitter, Jeffrey S., "Design and Analysis of Dynamic Huffman Codes", *Journal of the Association for Computing Machinery*, vol. 34, No. 4, (1987), pp. 825-845.

Wallace, Gregory K., "Overview of the JPEG (ISO/CCITT) Still Image Compression Standard", *Joint Photographic Expert Group Conference*, Crowborough; JPEG Forum Ltd., GB, XP017203197, (Nov. 1, 1989), pp. 1-14.

Wiegand, Thomas et al., "Overview of the H.264/AVC Video Coding Standard", *IEEE Transactions on Circuits and Systems for Video Technology*, 13(7), (Jul. 2003), pp. 560-576.

Xiao, Feng "DCT-Based Video Quality Evaluation", *Final Project for EE392J*, Winter 2000, (2000), 10 pages.

"Foreign Office Action", CN Application No. 201180010045.1, Nov. 2, 2015, 7 pages.

"Foreign Office Action", CN Application No. 201180010061.0, Nov. 11, 2015, 11 pages.

"Foreign Office Action", CN Application No. 201180010082.2, Dec. 9, 2015, 8 pages.

"Foreign Office Action", CN Application No. 201180010289.X, Nov. 2, 2015, 11 pages.

"Foreign Office Action", CN Application No. 201180010098.3, Oct. 21, 2015, 7 pages.

"International Preliminary Report on Patentability", Application No. PCT/EP2011/052446, Aug. 30, 2012, 8 pages.

"International Preliminary Report on Patentability", Application No. PCT/EP2011/052463, Aug. 30, 2012, 8 pages.

* cited by examiner

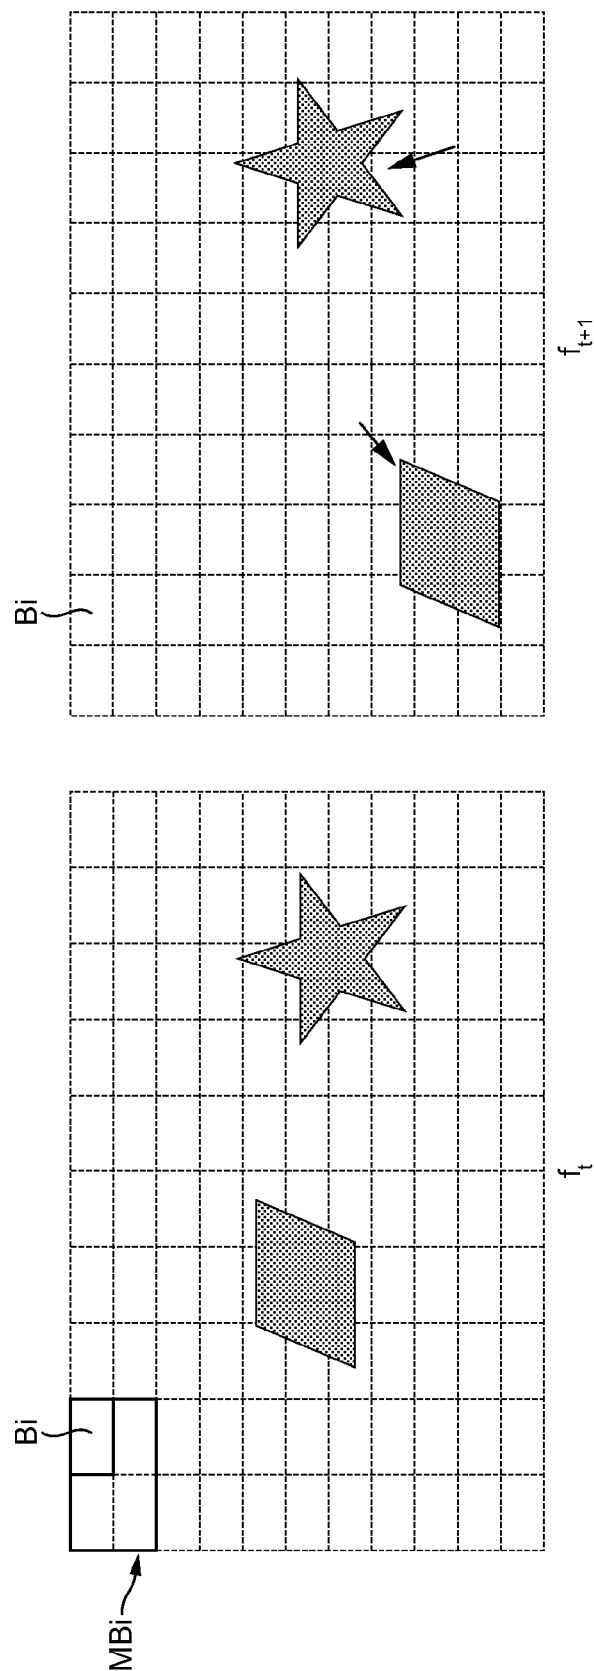


FIG. 1

+-----+-----+
130, 129, 129, 127, 125, 124, 124, 127,	129, 126, 123, 117, 109, 104, 101, 99,
128, 128, 127, 126, 126, 124, 125, 126,	128, 125, 122, 118, 110, 105, 102, 100,
126, 126, 126, 126, 125, 126, 126,	127, 124, 122, 119, 114, 108, 105, 103,
125, 125, 125, 125, 125, 125, 127,	126, 125, 123, 120, 115, 111, 107, 105,
124, 125, 125, 125, 125, 125, 126,	125, 124, 123, 122, 119, 115, 112, 109,
124, 125, 125, 125, 125, 125, 126,	125, 125, 124, 123, 120, 117, 114, 111,
122, 124, 124, 124, 124, 124, 125,	126, 126, 125, 124, 121, 119, 116, 114,
120, 122, 124, 124, 124, 124, 124,	126, 126, 126, 125, 123, 121, 118, 115,
+-----+-----+	
116, 121, 122, 123, 123, 123, 125, 127,	127, 127, 126, 125, 123, 120, 118, 115,
114, 120, 122, 121, 121, 121, 123, 125,	125, 126, 126, 124, 119, 118, 116, 115,
112, 118, 120, 119, 119, 120, 121,	122, 121, 120, 117, 115, 115, 114, 114,
111, 116, 118, 118, 118, 119, 120,	119, 117, 115, 113, 112, 112, 112, 112,
110, 114, 115, 116, 116, 116, 116,	115, 113, 111, 109, 109, 109, 109,
112, 114, 114, 115, 115, 115, 114, 114,	113, 111, 109, 108, 108, 107, 108,
113, 113, 113, 113, 113, 112, 112,	111, 110, 108, 106, 106, 106, 109,
114, 114, 114, 114, 114, 113, 113,	113, 111, 110, 109, 109, 109, 101,
+-----+-----+

FIG. 1a

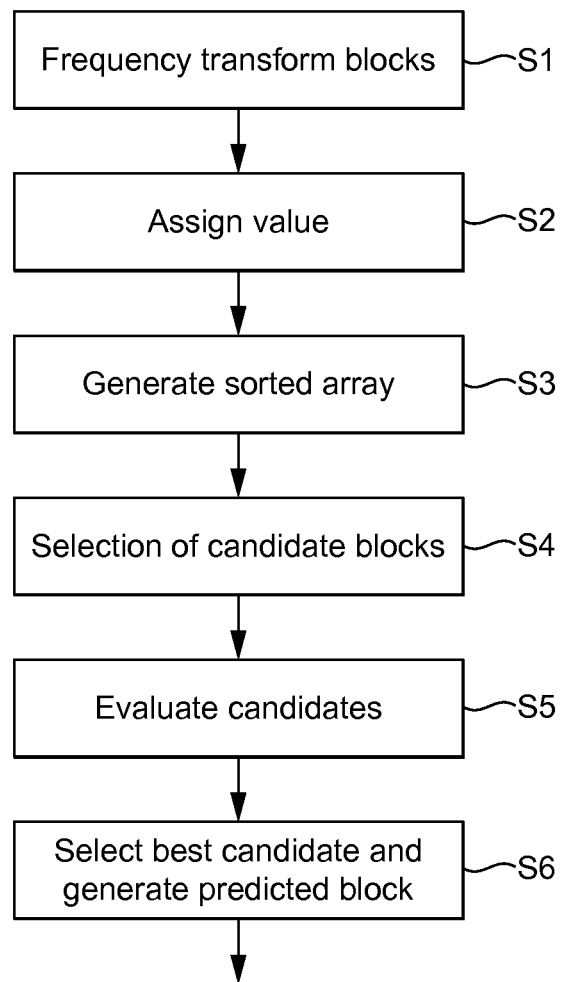



FIG. 2

Initial Blocks															
-320	304	0	0	0	0	0	0	-320	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	-304	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	400	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	304	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	-640	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	-304	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	400	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
-320	304	0	0	0	0	0	0	-640	0	0	0	0	0	0	0
304-288	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	-400	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	304	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	-320	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	320-304	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	-304	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
320	0	0	0	0	0	0	0	-320-304	0	0	0	0	0	0	0
-304	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	400	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

FIG. 3a



Initial Blocks															
-320	304	0	0	0	0	0	0	0	-320	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	-304	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	400	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	304	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	-640	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	-304	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	400	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
-320	304	0	0	0	0	0	0	0	-640	0	0	0	0	0	0
304	-288	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	-400	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	304	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	-320	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	320	-304	0	0	0	0	0
0	0	0	0	0	0	0	0	0	-304	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
320	0	0	0	0	0	0	0	0	-320	-304	0	0	0	0	0
-304	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	400	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	2	3	4	5	6	7	8	9	Bins Before Prediction						
0	-288	304	-304	320	-320	400	-400	-640	Values in each Bin						
487	1	5	6	2	5	3	1	2	Frequency						
1	2	3	4	5	6	7	8	9	#Bins After Prediction						
0	-288	304	-304	320	-320	400	-400	-640	Values in each Bin						
498	1	2	3	1	3	2	1	1	Frequency						
1	2	3	4	5	6	7	8	9	Values in each Bin						
0	-304	304	-320	400	320	-640	-288	-400	Frequency						
487	6	5	5	3	2	2	1	1							
	2.06	2.32	2.32	3.06	3.64	3.64	4.64	4.64							
	12.35	11.61	11.61	9.18	7.29	7.29	4.64	4.64	68.61 bits	Total Bits Before Prediction					
1	2	3	4	5	6	7	8	9							
0	-320	-304	400	304	-288	320	-400	-640							
498	3	3	2	2	1	1	1	1	Frequency						
	2.22	2.22	2.81	2.81	3.81	3.81	3.81	3.81							
	6.67	6.67	5.61	5.61	3.81	3.81	3.81	3.81	39.79 bits	Total Bits Before Prediction					




FIG. 3b Cont'd




FIG. 3

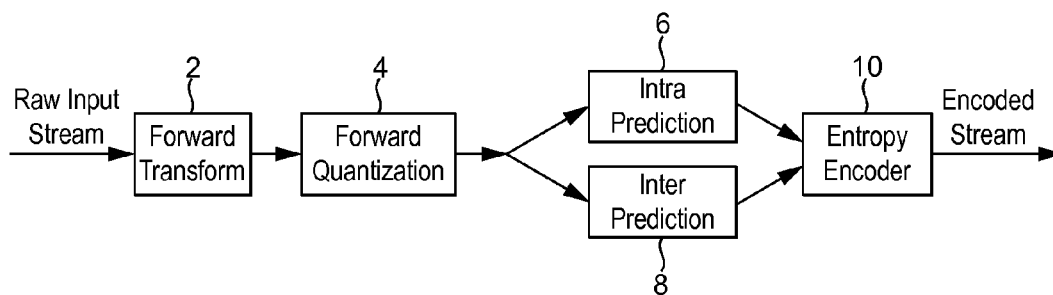


FIG. 4

b0				b1				b2				b3			
2464	-4032	-2688	-1728	224	576	-224	576	2688	0	224	-576	896	-576	-672	0
864	-1104	288	0	0	-368	288	-368	1728	-1840	0	0	864	-1104	288	0
672	-576	-448	288	224	-576	0	-288	448	-288	-448	288	672	0	224	0
0	0	0	0	0	0	0	0	-288	368	0	0	-288	0	0	0
-1344	-1440	1792	864	448	288	-224	0	224	0	0	0	-224	0	0	0
0	1104	-576	-368	576	736	-288	0	288	0	0	0	0	0	0	0
448	0	-448	0	224	-576	-224	0	224	0	0	0	0	0	0	0
288	0	0	0	0	-368	0	0	0	0	0	0	0	0	0	0
224	0	0	0	0	0	0	0	-672	576	0	0	448	288	-224	0
0	0	0	0	0	0	0	0	864	-736	0	0	-576	-368	288	0
0	0	0	0	0	0	0	0	-448	576	0	0	448	288	0	0
0	0	0	0	0	0	0	0	288	-368	0	0	-288	0	0	0
-672	864	-224	-288	0	0	0	0	-2016	0	-672	0	448	-288	-224	0
864	-1104	288	0	1152	-736	288	0	0	1104	0	736	-288	368	0	0
-448	288	0	0	0	0	-224	0	448	0	0	0	672	576	0	0
288	0	0	0	-288	-368	0	0	-288	736	0	0	0	0	0	0
b4				b8				b12				b16			
336	0	0	0	336	0	0	0	b17				9 coeff less out of 132 bin 122 is removed bin 1008 is removed bin 224 increased bin 288 increased			
288	0	0	0	576	0	0	0	b19							
-224	0	0	0	-224	0	0	0	b21							
0	0	0	0	-288	0	0	0	b23							
b18				b20				b22							
-336	288	0	0	-784	0	0	0	b18				b16			
288	0	0	0	0	0	0	0	b18				b16			
0	0	0	0	0	0	0	0	b18				b16			
-1008	288	224	288	-1232	0	0	0	b18				b16			
-864	368	0	0	-1152	0	0	0	b18				b16			
448	-288	0	0	672	0	0	0	b18				b16			
576	0	0	0	576	0	0	0	b18				b16			
112	0	0	0	336	0	0	0	b18				b16			
0	0	0	0	0	0	0	0	b18				b16			
0	0	0	0	0	0	0	0	b18				b16			

FIG 5a

FIG. 5a




FIG. 5a Cont'd

Frame 1
MB 1

-448	576	0	0	-448	-288	224	0	0	0	0	0	224	-576	0	0
288	0	288	0	288	0	-288	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	-864	0	288	-224	288	0	0	0	0	0	0	0	0	0	0
-288	736	-576	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	224	-288	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
-672	288	0	0	672	0	-448	0	-448	-288	224	0	-224	0	0	0
0	0	0	0	-288	736	576	0	288	-368	0	0	0	0	0	0
0	0	0	0	0	-288	0	288	0	0	0	0	0	0	0	0
0	0	0	0	0	0	-288	0	0	0	0	0	0	0	0	0
224	0	0	0	-224	-576	0	0	-224	-288	224	0	-224	0	0	0
0	0	0	0	576	-368	-288	0	-288	736	-288	0	288	0	0	0
0	0	0	0	0	0	0	0	-448	-288	224	0	0	0	0	0
0	0	0	0	0	0	0	0	288	-368	-288	0	0	0	0	0
224	288	224	0	0	0	0	0	0	0	0	0	0	0	0	0
-228	-368	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
672	288	-224	0	0	288	0	0	0	0	0	0	0	0	0	0
-228	368	0	0	-288	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
336	576	224	0	112	0	0	0	0	0	0	0	0	0	0	0
-288	-368	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2128	0	-224	0	1008	0	0	0	0	0	0	0	0	0	0	0
0	736	0	0	-288	-368	0	0	0	0	0	0	0	0	0	0
-224	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Frame 2
MB 0

2464	-4032	-2688	-1728	224	576	-224	576	2688	0	224	-576	896	-576	-672	0
864	-1104	288	0	0	-368	288	-368	1728	-1840	0	0	864	-1104	288	0
672	-576	-448	288	224	-576	0	-288	448	-288	448	288	672	0	224	0
0	0	0	0	0	0	0	0	-288	368	0	0	-288	0	0	0
-1344	-1440	1792	864	448	288	-224	0	224	0	0	0	-224	0	0	0
0	1104	-576	-368	576	736	-288	0	288	0	0	0	0	0	0	0
448	0	-448	0	224	0	-224	0	224	0	0	0	0	0	0	0
288	0	0	0	0	-368	0	0	0	0	0	0	0	0	0	0
224	0	0	0	0	0	0	0	-672	576	0	0	448	288	-224	0
0	0	0	0	0	0	0	0	864	-736	0	0	-576	-368	288	0
0	0	0	0	0	0	0	0	-448	576	0	0	448	288	0	0
0	0	0	0	0	0	0	0	288	-368	0	0	-288	0	0	0
-672	864	-224	-288	0	0	0	0	-2016	0	-672	0	448	-288	-224	0
864	-1104	288	0	1152	-736	288	0	0	1104	0	736	-288	368	0	0
-448	288	0	0	0	0	-224	0	448	0	0	0	672	576	0	0
288	0	0	0	-288	-368	0	0	-288	736	0	0	0	0	0	0
336	0	0	0	336	0	0	0	0	0	0	0	0	0	0	0
288	0	0	0	576	0	0	0	0	0	0	0	0	0	0	0
-224	0	0	0	-224	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	-288	0	0	0	0	0	0	0	0	0	0	0
-336	288	0	0	-784	0	0	0	0	0	0	0	0	0	0	0
288	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
-1008	288	224	288	-1232	0	0	0	0	0	0	0	0	0	0	0
-864	368	0	0	-1152	0	0	0	0	0	0	0	0	0	0	0
448	-288	0	0	672	0	0	0	0	0	0	0	0	0	0	0
576	0	0	0	576	0	0	0	0	0	0	0	0	0	0	0
112	0	0	0	336	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

FIG. 5b

Transmitted Coeff for Fr1 MB 1 Block 14

0	-288	224	0
0	736	-288	0
-224	-288	224	0
288	-368	-288	0

2 Coeff Less

Transmitted Coeff for Fr1 MB 1 Block 20

0	864	224	0
0	-368	0	0
0	0	0	0
0	0	0	0

2 Coeff Less

Transmitted Coeff for Fr2 MB 0 Block 0

2464	-4896	-2688	-1440
576	-368	-288	0
672	-576	-224	0
0	0	0	0

1 Coeff Less

Transmitted Coeff for Fr2 MB 0 Block 11

224	0	-448	0
-288	0	288	0
448	288	0	0
-288	0	0	0

1 Coeff Less

FIG. 5b Cont'd

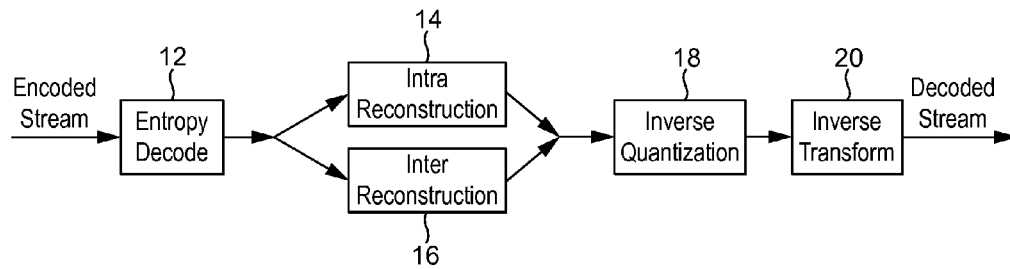


FIG. 6

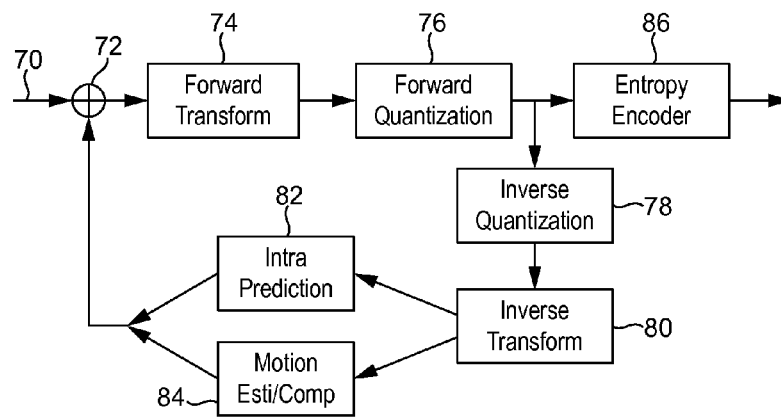


FIG. 7
(Prior Art)

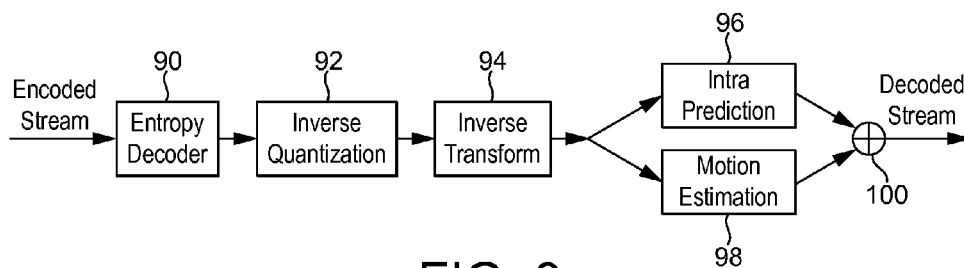


FIG. 8
(Prior Art)




FIG. 9

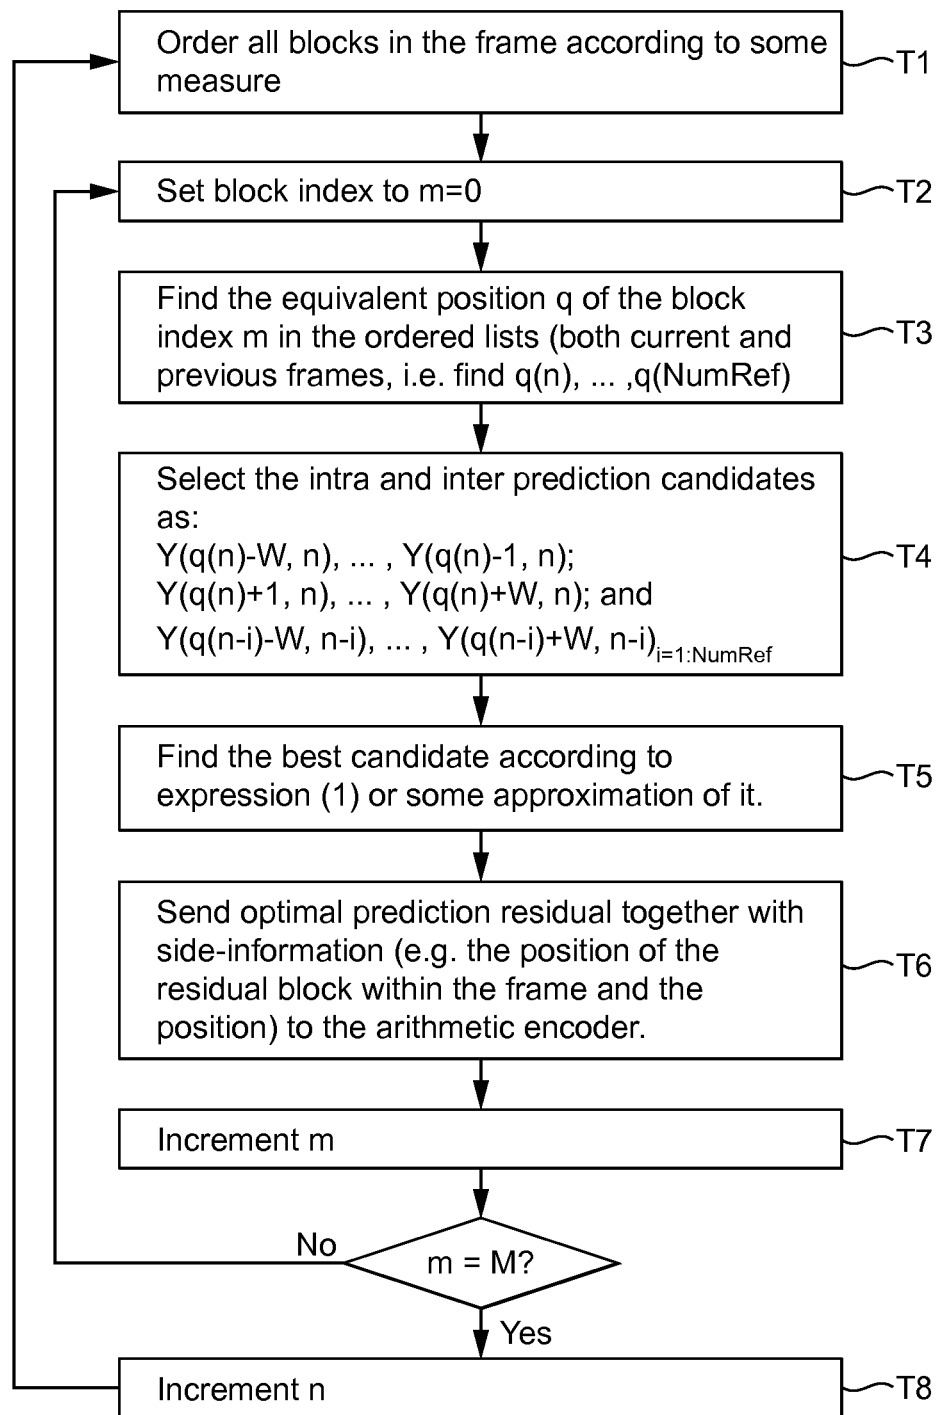


FIG. 10

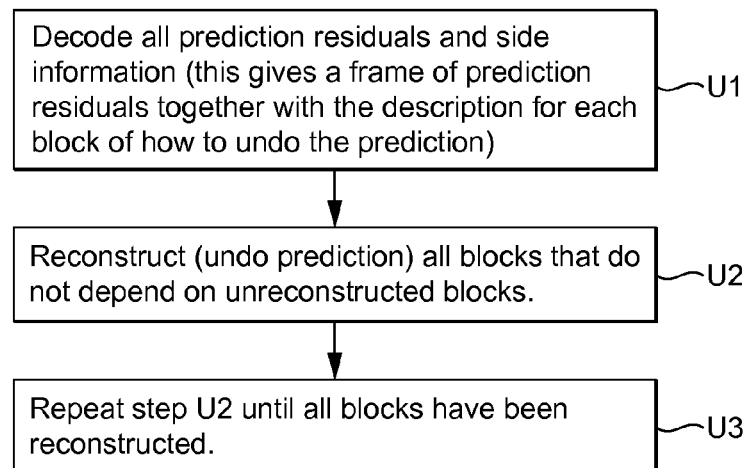


FIG. 11




FIG. 12




FIG. 13a




FIG. 13b




FIG. 13c





FIG. 13d






FIG. 14a




FIG. 14b




FIG. 14c




FIG. 15a




FIG. 15b




FIG. 15c




FIG. 15d




FIG. 15e




FIG. 15f




FIG. 15g




FIG. 15h

DATA COMPRESSION FOR VIDEO

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application No. 61/306,385, filed on Feb. 19, 2010 and also claims the benefit of U.S. Provisional Application No. 61/365,242, filed Jul. 16, 2010. The entire teachings of the above applications are incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates to the encoding and transmission of video streams.

BACKGROUND

In the transmission of video streams, efforts are continually being made to reduce the amount of data that needs to be transmitted whilst still allowing the moving images to be adequately recreated at the receiving end of the transmission. A video encoder receives an input video stream comprising a sequence of "raw" video frames to be encoded, each representing an image at a respective moment in time. The encoder then encodes each input frame into one of two types of encoded frame: either an intra frame (also known as a key frame), or an inter frame. The purpose of the encoding is to compress the video data so as to incur fewer bits when transmitted over a transmission medium or stored on a storage medium.

An intra frame is compressed using data only from the current video frame being encoded, typically using intra frame prediction coding whereby one image portion within the frame is encoded and signalled relative to another image portion within that same frame. This is similar to static image coding. An inter frame on the other hand is compressed using knowledge of a preceding frame (a reference frame) and allows for transmission of only the differences between that reference frame and the current frame which follows it in time. This allows for much more efficient compression, particularly when the scene has relatively few changes. Inter frame prediction typically uses motion estimation to encode and signal the video in terms of motion vectors describing the movement of image portions between frames, and then motion compensation to predict that motion at the receiver based on the signalled vectors. Various international standards for video communications such as MPEG 1, 2 & 4, and H.261, H.263 & H.264 employ motion estimation and compensation based on regular block based partitions of source frames. Depending on the resolution, frame rate, bit rate and scene, an intra frame can be up to 20 to 100 times larger than an inter frame. On the other hand, an inter frame imposes a dependency relation to previous inter frames up to the most recent intra frame. If any of the frames are missing, decoding the current inter frame may result in errors and artefacts.

These techniques are used for example in the H.264/AVC standard (see T. Wiegand, G. J. Sullivan, G. Bjontegaard, A. Luthra: "Overview of the H.264/AVC video coding standard," in IEEE Transactions on Circuits and Systems for Video Technology, Volume: 13, Issue: 7, page(s): 560-576, July 2003).

FIG. 7 illustrates a known video encoder for encoding a video stream into a stream of inter frames and interleaved intra frames, e.g. in accordance with the basic coding structure of H.264/AVC. The encoder receives an input video stream comprising a sequence of frames to be encoded (each divided into constituent macroblocks and subdivided into

blocks), and outputs quantized transform coefficients and motion data which can then be transmitted to the decoder. The encoder comprises an input **70** for receiving an input macroblock of a video image, a subtraction stage **72**, a forward transform stage **74**, a forward quantization stage **76**, an inverse quantization stage **78**, an inverse transform stage **80**, an intra frame prediction coding stage **82**, a motion estimation & compensation stage **84**, and an entropy encoder **86**.

The subtraction stage **72** is arranged to receive the input signal comprising a series of input macroblocks, each corresponding to a portion of a frame. From each, the subtraction stage **72** subtracts a prediction of that macroblock so as to generate a residual signal (also sometimes referred to as the prediction error). In the case of intra prediction, the prediction of the block is supplied from the intra prediction stage **82** based on one or more neighbouring regions of the same frame (after feedback via the reverse quantization stage **78** and reverse transform stage **80**). In the case of inter prediction, the prediction of the block is provided from the motion estimation & compensation stage **84** based on a selected region of a preceding frame (again after feedback via the reverse quantization stage **78** and reverse transform stage **80**). For motion estimation the selected region is identified by means of a motion vector describing the offset between the position of the selected region in the preceding frame and the macroblock being encoded in the current frame.

The forward transform stage **74** then transforms the residuals of the blocks from a spatial domain representation into a transform domain representation, e.g. by means of a discrete cosine transform (DCT). That is to say, it transforms each residual block from a set of pixel values at different Cartesian x and y coordinates to a set of coefficients representing different spatial frequency terms with different wavenumbers k_x and k_y (having dimensions of $1/\text{wavelength}$). The forward quantization stage **76** then quantizes the transform coefficients, and outputs quantised and transformed coefficients of the residual signal to be encoded into the video stream via the entropy encoder **86**, to thus form part of the encoded video signal for transmission to one or more recipient terminals.

Furthermore, the output of the forward quantization stage **76** is also fed back via the inverse quantization stage **78** and inverse transform stage **80**. The inverse transform stage **80** transforms the residual coefficients from the frequency domain back into spatial domain values where they are supplied to the intra prediction stage **82** (for intra frames) or the motion estimation & compensation stage **84** (for inter frames). These stages use the reverse transformed and reverse quantized residual signal along with knowledge of the input video stream in order to produce local predictions of the intra and inter frames (including the distorting effect of having been forward and reverse transformed and quantized as would be seen at the decoder). This local prediction is fed back to the subtraction stage **72** which produces the residual signal representing the difference between the input signal and the output of either the local intra frame prediction stage **82** or the local motion estimation & compensation stage **84**. After transformation, the forward quantization stage **76** quantizes this residual signal, thus generating the quantized, transformed residual coefficients for output to the entropy encoder **86**. The motion estimation stage **84** also outputs the motion vectors via the entropy encoder **86** for inclusion in the encoded bitstream.

When performing intra frame encoding, the idea is to only encode and transmit a measure of how a portion of image data within a frame differs from another portion within that same frame. That portion can then be predicted at the decoder (given some absolute data to begin with), and so it is only

necessary to transmit the difference between the prediction and the actual data rather than the actual data itself. The difference signal is typically smaller in magnitude, so takes fewer bits to encode.

In the case of inter frame encoding, the motion compensation stage **84** is switched into the feedback path in place of the intra frame prediction stage **82**, and a feedback loop is thus created between blocks of one frame and another in order to encode the inter frame relative to those of a preceding frame. This typically takes even fewer bits to encode than an intra frame.

FIG. **8** illustrates a corresponding decoder which comprises an entropy decoder **90** for receiving the encoded video stream into a recipient terminal, an inverse quantization stage **92**, an inverse transform stage **94**, an intra prediction stage **96** and a motion compensation stage **98**. The outputs of the intra prediction stage and the motion compensation stage are summed at a summing stage **100**.

There are many known motion estimation techniques. Generally they rely on comparison of a block with one or more other image portions from a preceding frame (the reference frame). Each block is predicted from an area of the same size and shape as the block, but offset by any number of pixels in the horizontal or vertical direction or even a fractional number of pixels. The identity of the area used is signalled as overhead ("side information") in the form of a motion vector. A good motion estimation technique has to balance the requirements of low complexity with high quality video images. It is also desirable that it does not require too much overhead information.

In the standard system described above, it will be noted that the intra prediction coding and inter prediction coding (motion estimation) are performed in the unquantized spatial domain.

More recently, motion estimation techniques operating in the transform domain have attracted attention. However, none of the existing techniques are able to perform with low complexity (thus reducing computational overhead), while also delivering high quality. Hence no frequency domain techniques for motion estimation are currently in practical use.

The VC-1 video codec has an intra prediction mode which operates in the frequency domain, in which the first column and/or first row of AC coefficients in the DCT (Discrete Fourier Transform) domain are predicted from the first column (or first row) of the DCT blocks located immediately to the left or above the processed block. That is to say, coefficients lying at the edge of one block are predicted from the direct spatial neighbours in an adjacent block. For reference see "The VC-1 and H.264 Video Compression Standards for Broadband Video Services", AvHari Kalva, Jae-Beom Lee, pp. 251.

SUMMARY OF INVENTION

The present invention provides a technique for performing one or more aspects of video coding such as quantization, intra prediction coding or inter prediction coding in dependence on a perceptual model taking into account human sensitivity to data in the video signal. The perceptual model may relate to spatial frequency, temporal frequency, contrast sensitivity, colour sensitivity, a structural metric, and/or one or more parameters affecting perception such as motion in the video, the distance of a recipient user from the screen, and the size, aspect ratio or resolution of the screen of the recipient terminal.

According to one aspect of the present invention, there may be provided a method of encoding a video signal for trans-

mission to a decoder of a receiving terminal, the method comprising: receiving a video signal comprising a plurality of frames, each frame comprising one or more portions of video data; quantizing the video data of each portion; and encoding the video data of each portion; wherein the quantization is performed according to a quantization bin size distribution which uses different quantization bin sizes for different video data, wherein the quantization bin size distribution is related to a measure of human sensitivity to the video data.

In embodiments, the method may comprise transforming each of said portions into a transform domain, so as for each portion to generate transformed video data comprising a set of coefficients, each coefficient representing a contribution to a property of the portion at a different spatial and/or temporal frequency; said quantization may comprise quantizing the coefficients of each portion; and the quantization bin size distribution may use different quantization bin sizes for different transformed video data.

The quantization bin size distribution may be related to a measure of human sensitivity to said property at different spatial and/or temporal frequencies.

The quantization bin size distribution may use different quantization bin sizes for different ones of said coefficients.

The property may comprise one of a luminance and a colour property.

Some of said coefficients may represent luminance and some of said coefficients may represent at least one colour property, and the method may comprise comprising using a different quantization bin size distribution for quantizing the luminance and colour.

Some of said coefficients may represent a first colour channel and some of said coefficients represent a second colour channel, and the method may comprise different quantization bin size distributions for the first and second colour channels.

The quantization bin size distribution may be related to a measure of human sensitivity to a visual structure of the video data.

The quantization bin size distribution may be made variable in dependence on a parameter which affects human perception.

The quantization bin size distribution may be dynamically varied during ongoing encoding in dependence on a parameter which affects human perception, so as to apply a different quantization bin size distribution for different ones of said frames or different portions of the frames.

The quantization bin size distribution may be varied dynamically in dependence on motion in the video signal, the motion being a parameter which affects human perception.

The motion may be determined by a motion estimation algorithm applied to the frames of the video signal.

The motion may be detected by physical sensors in a mobile terminal housing a camera which generates said video signal.

The quantization bin size distribution may be controlled in dependence on a distance of a user of the receiving terminal from a screen of the receiving terminal, said distance being a parameter which affects human perception.

The method may comprise receiving an indication of said distance from the receiving terminal for using in controlling the quantization bin size distribution.

Said distance may be determined by an autofocus of a camera of said receiving terminal.

Said distance may be determined by a user setting.

Said bin size distribution may be controlled in dependence on at least one of an aspect ratio, a size and a resolution of a screen of said recipient terminal, the aspect ratio, size and/or resolution being a parameter which affects human perception.

5

The method may comprise receiving an indication of said aspect ratio, screen size and/or resolution from the receiving terminal for using in controlling the quantization bin size distribution.

The method may comprise transmitting an indication of the bin size distribution to the decoder for use in decoding the video signal.

The video data may be input to the quantizer before inter or intra encoding such that the video data represents absolute values.

According to another aspect of the present invention, there may be provided a method of compressing a video signal for transmission to a decoder of a receiving terminal, the method comprising: receiving a video signal comprising a plurality of frames, each frame comprising a plurality of portions of video data;

for each of a plurality of target ones of said portions, generating difference data representing a difference between the target image portion and a respective reference portion; and outputting the difference data in an encoded bitstream; wherein the method comprises, prior to generating the difference data, modifying at least one of said target and reference portions based on a measure of human sensitivity to the video data, so as to reduce a number of bits required to encode the difference data.

In embodiments, the modification may reduce a magnitude of the difference data and thereby reduce the number of bits required to encode the difference data.

The method comprise outputting the difference data through an entropy encoder, and said modification may reduce a variation in said difference data so as to reduce entropy and thereby reduce the number of bits required to encode the difference data.

The method may comprise transforming each of said portions into a transform domain, so as for each portion to generate transformed video data comprising a set of transform domain coefficients, each coefficient representing a contribution to a property of the portion at a different spatial and/or temporal frequency; and for each of a plurality of target ones of said portions, the difference data may represent a difference between the coefficients of a target image portion and transform domain coefficients of the respective reference portion.

Said modification may comprise modifying one or more of the transform domain coefficients of at least one of said target and reference portions based on a measure of human sensitivity to the transform domain coefficients, thereby reducing the number of bits required to encode the difference data.

The property may comprise one of a luminance and a colour property.

The modification may be based on a measure of human sensitivity to a visual structure of the video data.

The degree of modification may be made variable in dependence on a parameter which affects human perception.

The degree of modification may be dynamically varied during ongoing encoding in dependence on a parameter which affects human perception.

The degree of modification may be varied dynamically in dependence on motion in the video signal, the motion being a parameter which affects human perception.

The motion may be determined by a motion estimation algorithm applied to the frames of the video signal.

The motion may be detected by physical sensors in a mobile terminal housing a camera which generates said video signal.

The degree of modification may be controlled in dependence on a distance of a user of the receiving terminal from a

6

screen of the receiving terminal, said distance being a parameter which affects human perception.

The method may comprise receiving an indication of said distance from the receiving terminal for using in controlling the degree of modification.

Said distance may be determined by an autofocus of a camera of said receiving terminal.

Said distance may be determined by a user setting.

The degree of modification may be controlled in dependence on at least one of an aspect ratio, a size and a resolution of a screen of said recipient terminal, the aspect ratio, size and/or resolution being a parameter which affects human perception.

The method may comprise receiving an indication of said aspect ratio, screen size and/or resolution from the receiving terminal for using in controlling the degree of modification.

For each target portion, the method may comprise selecting the respective reference portion from amongst a plurality of potential reference portions.

The plurality of potential reference portions may comprise one or more portions from the same frame as the target portion, the method thus being usable for intra encoding.

The plurality of potential reference portions may comprise one or more portions from a different frame as the target portion, the method thus being usable for intra encoding.

The respective reference portion may be selected based on the number of bits that would be required to encode the difference data.

The method may comprise determining an overall energy of each of the target portion and the reference portions, and determining a subset of candidate portions each having a block energy within a threshold range of the target portion, wherein the respective reference portion may be selected from amongst the candidate portions.

According to another aspect of the present invention, there may be provided a method of transmitting an encoded bitstream comprising: encoding the video signal according to any of the above method features, and transmitting the encoded bit stream to the decoder at the receiving terminal.

According to another aspect of the present invention, there is provided a method of compressing video data comprising: providing for an image portion coefficients defining image data of the image portion; arranging the coefficients into sets and determining an index value for each set based on objective and perceptual metrics for the set; sorting the sets based on the index values to generate a sorted list; selecting at least two matching sets based on proximity in the list and perceptual metrics; and generating residuals for one of said matching sets based on the comparison of coefficients of that one with the other of the matching sets.

A further aspect provides a method as defined above, wherein the step of selecting matching sets comprises determining a group of candidates based on proximity in the list, and selecting a best match from the candidates, wherein the best match is based on bit rate contribution and perceptual quality.

According to another aspect of the present invention, there is provided a method of compressing video data comprising: receiving a stream of video data defining a video image; for each of a plurality of image portions defined by temporal and/or spatial data, transforming the temporal and/or spatial data into the frequency domain to generate for each image portion a plurality of sets of frequency domain coefficients; determining an index value for each set based on objective and perceptual metrics for the set; selecting at least two matching sets of frequency domain coefficients using the

index values; and generating residuals for one of the matching sets based on comparison of that one with the other of the matching sets.

In embodiments the image data may be chrominance and luminance data. Each set may represent a block or object. The image portion may comprise a macroblock having a two-dimensional array of pixels. The image portion may comprise a single frame of a video stream, wherein the method may be utilised to compress video data in an intra frame. The image portion may comprise a plurality of frames, wherein the method may be used for inter frame prediction. Each set may comprise an array of coefficients.

The index value may represent energy of each set. The index value may be based on the number of coefficients having a value of zero. The index value may be an aggregate combination of coefficients. The index value may be the average of the moduli of the coefficients in the set.

The step of selecting matching sets may comprise determining a group of candidates based on proximity in the list, and selecting a best match from the candidates by comparing bit rate contribution and perceptual quality.

According to another aspect of the present invention, there may be provided a method of transmitting video data from a transmitting terminal to a recipient terminal, the method comprising: generating a stream of video data representing a moving image; compressing the video data in accordance with a method of any of the above method features; and transmitting coefficients from the video stream for sets which are not predicted and the residuals for the at least one set which is predicted as part of an encoded video stream.

According to another aspect of the invention, there may be provided an encoder comprising signal processing apparatus configured to operate in accordance with any of the above method features.

According to another aspect of the invention, there may be provided a computer program product comprising software embodied on a computer-readable medium and configured so as when executed on a processor to perform operations in accordance with any of the above method features.

According to another aspect of the invention, there may be provided a bitstream encoded according to any of the above method features. According to another aspect, there may be provided a network equipment comprising a transmission medium carrying the bitstream.

According to another aspect of the invention, there may be provided a method of decoding a video signal encoded according to any of the above method features.

The method may comprise decoding the encoded bitstream based on an indication of the quantization bin size distribution received in the encoded bitstream.

According to another aspect of the present invention, there may be provided a corresponding decoder comprising signal processing apparatus configured to operate in accordance with the method of decoding.

According to another aspect of the invention, there may be provided a decoder program product comprising software embodied on a computer-readable medium and configured so as when executed to perform operations in accordance with the method of decoding.

BRIEF DESCRIPTION OF THE DRAWINGS

For a better understanding of the present invention and to show how the same may be carried into effect, reference will now be made by way of example, to the accompanying drawings, in which:

FIG. 1 is a schematic illustration of two frames of a video stream,

FIG. 1a shows the pixel values of blocks represented in the spatial domain,

FIG. 1b shows coefficients of blocks represented in the frequency domain,

FIG. 2 is a flow diagram showing an encoding method,

FIG. 3 is a graph showing a sorted block list,

FIG. 3a illustrates an example of block-sorting,

FIG. 3b illustrates an example of block matching prediction

FIG. 4 is a schematic block diagram of an encoder,

FIG. 5A is a schematic example of an intra frame prediction,

FIG. 5B is a schematic example of an inter frame prediction,

FIG. 6 is a schematic diagram of a decoder,

FIG. 7 is a schematic block diagram of an encoder,

FIG. 8 is a schematic block diagram of a decoder,

FIG. 9 is a schematic illustration of selection of candidates for block matching,

FIG. 10 is a flow chart of an encoding method,

FIG. 11 is a flow chart of a decoding method,

FIG. 12 is a schematic representation of a transmitted bitstream,

FIG. 13a is a schematic illustration of block scaling,

FIG. 13b is a schematic illustration of block rotation,

FIG. 13c is another schematic illustration of block rotation,

FIG. 13d is another schematic illustration of block rotation,

FIG. 13e is a diagram showing a calculation for use in block rotation,

FIG. 14a schematically illustrates a motion shift between two frames,

FIG. 14b is another schematic illustration of a motion shift,

FIG. 14c schematically shows using a motion shift to reduce data transmission,

FIGS. 15a-15h schematically represent aspects of a perceptual model.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

In the following there is described a method and system for data compression in a video transmission system. First is described an exemplary technique of block matching performed in the frequency domain for selecting a reference block to use in prediction coding of a target block. Next are described some exemplary techniques for signalling the identity of reference blocks for use in video prediction coding, and a scheme for further reducing the bitrate of a transmitted video stream. Following that are described some exemplary image processing techniques that can be advantageously performed in the frequency domain, a scheme for further reducing the bitrate of a transmitted video stream, and a perceptual model for use in compressing video data. In particularly preferred embodiments these techniques may be combined, but alternatively they may be used independently.

Block Matching

FIG. 1 schematically illustrates two successive frames f_t and f_{t+1} of a video image at two respective moments in time t and $t+1$. For the purpose of inter frame prediction the first frame f_t may be considered a reference frame, i.e. a frame which has just been encoded from a moving sequence at the encoder, or a frame which has just been decoded at the decoder. The second frame f_{t+1} may be considered a target frame, i.e. the current frame whose motion is sought to be

estimated for the purpose of encoding or decoding. An example with two moving objects is shown for the sake of illustration.

Motion estimation is itself known in the art and so is described herein only to the extent necessary to provide suitable background for the present invention. According to International Standards for Video Communications such as MPEG 1, 2 & 4 and H.261, H.263 & H.264, motion estimation is based on block-based partitions of source frames. For example, each block may comprise an array of 4×4 pixels, or 4×8, 8×4, 8×8, 16×8, 8×16 or 16×16 in various other standards. An exemplary block is denoted by B_i in FIG. 1. The number of pixels per block can be selected in accordance with the required accuracy and decode rates. Each pixel can be represented in a number of different ways depending on the protocol adopted in accordance with the standards. In the example herein, each pixel is represented by chrominance (U and V) and luminance (Y) values (though other possible colour-space representations are also known in the art). In this particular example chrominance values are shared by four pixels in a block. A macroblock MB_i typically comprises four blocks, e.g. an array of 8×8 pixels for 4×4 blocks or an array of 16×16 pixels for 8×8 blocks. Each pixel has an associated bit rate which is the amount of data needed to transmit information about that pixel.

FIG. 2 is a schematic flow chart of a data compression method in accordance with a preferred embodiment of the present invention. The method preferably uses block matching based on objective metrics. That is, one or more metrics of a current target block to be encoded are compared to the corresponding metrics of a plurality of other blocks, and a reference block is selected based on a measure of similarity of those metrics. The reference block then forms the basis for encoding the current block by means of prediction coding, either intra-frame coding in the case where the reference block is from the same frame f_{t+1} or inter-frame coding where the reference block is from a preceding frame f_t (or indeed f_{t-1} , or f_{t-2} , etc.). The idea behind the block matching is to choose a reference block which will result in a small residual signal when the current block is encoded relative to that reference block (i.e. so that the difference between the actual current block and the prediction will be small when predicted from the selected reference block), hence requiring only a small number of bits to encode.

It is a particularly preferred aspect of the technique that block matching is carried out in the frequency domain, i.e. based on comparison of one or more metrics of a transformed representation of the blocks.

Hence at step S1, a frequency domain transform is performed on each portion of the image of each of a plurality of frames, e.g. on each block. Each block is initially expressed as a spatial domain representation whereby the chrominance and luminance of the block are represented as functions of spatial x and y coordinates, $U(x,y)$, $V(x,y)$ and $Y(x,y)$ (or other suitable colour-space representation). That is, each block is represented by a set of pixel values at different spatial x and y coordinates. A mathematical transform is then applied to each block to transform into a transform domain representation whereby the chrominance and luminance of the block (or such like) are represented as a function of variables such as wavenumbers k_x and k_y having dimensions of 1/wavelength, i.e. $U(k_x, k_y)$, $V(k_x, k_y)$ and $Y(k_x, k_y)$. That is, the block is transformed to a set of coefficients representing the amplitudes of different spatial frequency terms which can be considered to make up the block. Possibilities for such transforms include the Discrete Cosine transform (DCT), Karhunen-Loeve Transform (KLT), or others. E.g. for a block of $N \times M$

pixels at discrete x and y coordinates within the block, a DCT would transform the luminance $Y(x,y)$ to a set of frequency domain coefficients $Y(k_x, k_y)$:

$$Y(k_x, k_y) = \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} Y(x, y) \cos \left[\frac{\pi k_x}{2N} (2x+1) \right] \cos \left[\frac{\pi k_y}{2M} (2y+1) \right]$$

And inversely, the x and y representation $Y(x,y)$ can be determined from a sum of the frequency domain terms summed over k_x and k_y . Hence each block can be represented as a sum of one or more different spatial frequency terms having respective amplitude coefficients $Y(k_x, k_y)$ (and similarly for U and V). The transform domain may be referred to as the frequency domain (in this case referring to spatial frequency).

In some embodiments of the invention, the transform could be applied in three dimensions. A short sequence of frames effectively form a three dimensional cube or cuboid $U(x,y,t)$, $V(x,y,t)$ and $Y(x,y,t)$. In the case of a three dimensional transform, these would transform to $U(k_x, k_y, f)$, $V(k_x, k_y, f)$ and $Y(k_x, k_y, f)$. The term “frequency domain” may be used herein may be used to refer to any transform domain representation in terms of spatial frequency (1/wavelength domain) transformed from a spatial domain and/or temporal frequency (1/time period domain) transformed from a temporal domain.

Once the blocks are transformed into the frequency domain, block matching is performed by comparing the transformed frequency domain coefficients of the current block to those of a plurality of other blocks. A reference block for prediction coding of the current block (either intra or inter) can then be selected based on a measure of block similarity determined from the frequency domain coefficients.

An advantage of block-matching in the frequency domain is that the transform tends to compact the energy of a block into only a few non-zero (or non-negligible) coefficients, and thus that comparison can now be based only on only a few frequency coefficients instead of all the coefficients in the block. That is, since the frequency transform concentrates the energy into only a few significant coefficients, then efficient block matching (or indeed other processing) can be performed by only considering those few significant coefficients. This technique thus provides a unique approach to the problem of data compression in video transmission. Although not every pixel need be directly compared when comparing patterns, nevertheless, a complete search can be achieved.

For example consider an illustrative case as shown in FIGS. 1a and 1b. Here, the representation of a block in the frequency domain is achieved through a transform which converts the spatial domain pixel values to spatial frequencies. FIG. 1a shows some example pixel values of four 8×8 blocks in the spatial domain, e.g. which may comprise the luminance values $Y(x, y)$ of individual pixels at the different pixel locations x and y within the block. FIG. 1b is the equivalent in the frequency domain after transform and quantization. E.g. in FIG. 1b such coefficients may represent the amplitudes $Y(k_x, k_y)$ of the different possible frequency domain terms that may appear in the sum. The size of the block in spatial and frequency domain is the same, i.e. in this case 8×8 values or coefficients. However, due to the properties of these transforms then the energy of the block is compacted into only few coefficients in the frequency domain, so the entire block can be considered by processing only these few coefficients.

As can be seen from this example, only four values need to be processed to find a match for these four blocks in the

11

frequency domain, whereas in the spatial domain there are 256 values that would need to be processed. Thus unlike prior techniques, the present invention may allow a full true search to be performed but without the need to “touch” every pixel in the block, i.e. without the need to process each individual pixel.

It will be appreciated that while blocks and macroblocks are referred to herein, the techniques can similarly be used on other portions definable in the image. Frequency domain separation in blocks and/or portions may be dependent on the choice of transform. In the case of block transforms, for example, like the Discrete Cosine transform (DCT) and Karhunen-Loeve Transform (KLT) and others, the target block or portions becomes an array of fixed or variable dimensions. Each array comprises a set of transformed quantized coefficients. E.g. in the more detailed example of FIG. 5A, each macroblock MB of 16×16 pixels may be represented in the frequency domain by 16 luminance blocks and 8 chrominance blocks; each block $b_0 \dots b_{23}$ having a 4×4 array of quantized coefficients.

According to another preferred aspect of the present invention, block matching may be performed within a sorted list based on an index value reflecting the relative importance of the block. In this case the selection of matching blocks may be performed based on an aggregate of values used for the importance indexing. A preferred example will now be described with reference to steps S2 to S6 of FIG. 2 and the example blocks of FIG. 5A.

At Step S2, each block $b_0 \dots b_{23}$ in the frequency domain is assigned an index value derived from one or more of its frequency domain coefficients. For example, the index value may represent the energy of the block. E.g. this may comprise an aggregate over the coefficients of the block, such as a number of zero coefficients, number of non-zero coefficients, or an average or total value of the moduli of the coefficients in each block.

At Step S3, the blocks from at least one frame are then sorted based on the index value. This may involve generating a sorted list in which the entries represent blocks ordered according to their index values, e.g. their block energies.

At Step S4, a subset of candidate blocks is identified from the sorted array by determining a search range or threshold Δ based on the index values. The candidate blocks will be potential matches as reference blocks for use in prediction coding of a current block to be encoded. This is illustrated in FIG. 3. For example this may be achieved by determining an energy range $\pm\Delta$ from the current block to be encoded, and determining that all blocks within that range of the current block are candidates for potential selection as a reference block (i.e. candidates for a “match” to the current block for the purpose of prediction coding).

At Step S5, the candidate blocks are then evaluated for similarity. For example, block similarity is preferably determined based on bit rate, where the bit rate is a measure of the number of bits that would need to be transmitted in order to define the residuals for the current block if predicted from each candidate block. An example of this will be discussed in more detail shortly.

At Step S6, the best matching candidate is determined based on its similarity, and the current target block is encoded relative to that matching candidate. The encoding comprises subtracting the frequency domain coefficients of the reference block from those of the current block in order to generate a residual signal, and then encoding the residual of the current block into the encoded bitstream along with the identity of the respective selected reference block (instead of encoding the target block’s actual absolute coefficients). The reference

12

block is thus used as a prediction of the current block. The residual is the difference between the frequency domain coefficients of the current block and the frequency domain coefficients of the reference block, which requires fewer bits to encode and so the encoding results in a compressed video signal. The best candidate for use as the reference block is preferably selected by calculating the bit rate that would be required to transmit the residuals for the current block based on the candidate plus overhead information identifying the candidate block, in comparison with the bit rate that would be required for other such candidates. It will be readily appreciated that a match does not imply identical blocks, but blocks that are sufficiently similar that residuals can be transmitted at a lower bit rate.

FIG. 3 is a graph illustrating the arrangement of a sorted array. The list of sorted blocks is shown on the horizontal axis, with block energy index value on the vertical axis. The block energy index value is an example of an objective metric derived from the block’s coefficients.

As described above, a best matching reference block is selected having an index within a certain search range or threshold Δ . Thus according to one preferred aspect, the invention provides a method of searching amongst the blocks for matches based on similarity of their indices. By searching for matches by their energy index or such like, this advantageously expands the potential for matches to anywhere within the frame or another frame. Hence the matching need not be restricted to adjacent regions of the target block. For instance, blocks having similar energies may achieve a good even if located on opposite sides of a frame, e.g. blocks of a similar background area appearing at different locations in the frame.

According to another preferred aspect of the invention, block matching is performed by first selecting a subset of candidate blocks based on a first metric (e.g. the index value), and then selecting a matching candidate block from within the subset based on a second metric (e.g. bitrate cost). The matching block is then used as a reference block in prediction coding of a current block to be encoded. One advantage of narrowing the possible matches down to a preliminary subset of candidates based on a first metric, particularly based on an aggregate metric such as block energy, is that unlikely candidates can be eliminated early on without incurring significant processing burden. That is, the sort may be used to discard unlikely candidates. Thus the more processor-intensive comparison based on the second metric, such as the bit rate comparison, need only be performed for a relatively small number of pre-vetted candidates, thus reducing the processing burden incurred by the block matching algorithm. E.g. blocks with very different block energies are unlikely to be good matches and therefore it is unlikely to be worth the processing cost of comparing their potential bitrate contributions. To minimize processing, the selection of a matching block in Step S6 is preferably performed within a small neighbourhood within the list (search range $\pm\Delta$).

Note though that the sort only gives a certain probability of a match and may be chosen depending on performance considerations. A smaller choice of Δ results in a lower processing cost but fewer candidates, risking not find the best possible match. A larger choice of Δ on the other hand incurs a higher processing cost but will include more candidates and so have a better chance of finding the best match. In embodiments, Δ could even be adapted dynamically based on one or more performance factors such as available up or downlink bandwidth or available processing resources. Note also that the same value of Δ need not necessarily be use in the $+\Delta$ direction as in the $-\Delta$ direction.

13

It will be appreciated that at Step S3, the sorted array can be generated for a macroblock (as shown in the example of FIG. 5A), for a single frame (for intra frame data compression) or for a current target frame and one or more reference frames (for inter frame motion estimation).

In one particularly advantageous embodiment, the same sorted list is used to match multiple target blocks (by determining respective subsets of candidates within the same list). Further, if the list contains blocks from both the current frame and one or more preceding frames, then the same list can even be used for both inter and intra matching within the same sorted list. E.g. when processing a particular target frame it may be that a good match may not be found within that frame, in which case the method may look to other frames since the complexity is low and the matching method is the same. According to preferred embodiments of the present invention, there is no need to use a different method for finding inter frame matches between frames than is used for intra matching within a frame.

By replacing an exhaustive, repetitive search performed for every block with a single sort that is performed once for an entire frame or even multiple frames, the selection of a matching block can be performed in a small neighbourhood using the sorted list. Preferably the sort is performed once for multiple frames, so that both inter and intra matches can be processed at the same stage over the same sorted list. E.g. this may involve looking for a match within the sorted list of the current frame and, if no satisfactory match is found, looking into the sorted lists of one or more other frames to find a better match.

The above-described aspect of the present invention thus provides a method of compressing video data which can be applicable both to intra frame compression and to inter frame motion estimation. In the past, algorithms have adopted different approaches to inter versus intra data compression. The invention on the other hand can advantageously provide a unified technique used for both intra and inter frame prediction.

Another benefit of the method is that due to its low complexity, the number of used reference frames can be substantially higher in comparison with existing algorithms.

Furthermore, note that conventional motion estimation predicts each block from an area offset by any arbitrary number of pixels or even fractional number of pixels in the horizontal or vertical direction, whereas the approach used in the present invention differs by restricting the prediction to performing only block-to-block matching. That is, matching on a block-by-block basis whereby a block is matched to another whole block (rather than any arbitrarily offset area requiring a motion vector specifying any number of pixels or fractional number of pixels). In a particularly advantageous combination of features, the block-to-block matching may be performed in the frequency domain where efficiency can be derived by predicting only a subset of frequency domain coefficients between two or more blocks.

Once a matching block has been selected at step S6 and the current target block has been encoded relative to that matching block, the residual of the frequency domain coefficients is output via an entropy encoder for inclusion in the encoded bitstream. In addition, side information is included in the bitstream in order to identify the reference block from which each encoded block is to be predicted at the decoder. Each block may be identified by its location, i.e. by its address or position within a particular frame. Each frame may be identified by a frame number. Because of the above distinction, note that the side information identifying the selected reference block may be signaled in the bitstream in the form of a

14

block address identifying the location of the reference block in terms of a whole number of blocks. This may take the form of an absolute block address, i.e. a position relative to a fixed point in the frame. Alternatively it may take the form of a relative address. The side information may also identify the frame of the selected reference block if candidates may be selected from a plurality of different potential frames.

This is different from a conventional motion vector, which is signaled in the form of a small vector relative to the current block, the vector being any number of pixels or fractional pixels.

As mentioned, the VC-1 video codec has an intra prediction mode in which the first column and/or first row of AC coefficients in the DCT domain are predicted from the first column (or first row) of the DCT blocks located immediately to the left or on the top of the processed block. However, this differs from the approach used in aspects of the present invention in that it is restricted to using only predetermined spatially-adjacent coefficients for intra prediction. VC-1 does not allow intra matching to a selected reference block, e.g. selected based on block energy and/or bitrate contribution (and therefore VC-1 also does involve signaling the identity of a selected reference block to the decoder).

FIG. 4 is a schematic block diagram showing the architecture of an encoding technique in accordance with one embodiment of the invention. The raw input video stream is received by a forward transform stage 2. The output of this stage is supplied to a forward quantization stage 4. The forward transform stage 2 applies spatial or spatial-temporal transform into the frequency domain as a first coding step. The forward quantization stage 2 applies quantization and generates for each block a set of quantized coefficients in the frequency domain. The transform coefficients from the forward quantization stage 2 of each intra frame in the temporal domain of the input video stream are supplied to an intra prediction stage 6.

The intra prediction stage 6 operates to locate candidate blocks for prediction within each frame, using the method described above. The transform coefficients of inter frames are supplied from the forward quantization stage 4 to an inter-prediction stage 8, which separates the candidate blocks for prediction of target frames as described above. The outputs of the intra prediction stage and the inter-prediction stage 8 are supplied to an entropy encoder 10 which encodes the data to provide an encoded stream for transmission. The encoded stream contains a sequence of information comprising, for each block, a set of coefficients (actual or residual), data defining whether the block is to be predicted and, if it is, an indication of the reference block from which it is to be predicted. The identity of the reference block may be represented in the encoded bitstream as an absolute block location within a frame, i.e. by reference to a fixed point, and not relative to the current block. Alternatively the location may be represented in the encoded bitstream as a difference between the location of the current block and the block from which it is predicted. Either way, the block location is expressed in terms of a number of intervals of whole blocks, i.e. as a block address, and so a benefit is achieved because this requires far less overhead to encode than a conventional motion vector expressing an offset in pixels or even fractions of pixels.

Note that the arrangement does not involve a loop back into the spatial domain as in the standard encoder of FIG. 7. Hence block matching is performed in the transformed frequency domain based on frequency domain coefficients.

Note also that in preferred embodiments, the selection of the reference block is performed in the quantized domain, i.e. a non-distorting, lossless environment. Therefore no addi-

15

tional distortion is applied to the candidate blocks or current blocks before performing the selection.

FIG. 5A illustrates schematically a prediction example. The case illustrated in FIG. 5A is where the technique is used for intra prediction between different blocks of the same macroblock in one frame. FIG. 5A illustrates on the left hand side luminance and chrominance data transformed into the frequency domain for a macroblock (16×16 pixels). The frequency transformed coefficients are organised into blocks **b0**, **b1**, etc, each block comprising a 4×4 array of coefficients. Blocks **b0** to **b15** represent luminance data (y) for the macroblock, and blocks **b16** to **b23** represent chrominance data (u,v) for the macroblock.

There are different schemes for treating the luma and chroma channels. A common way is the 4:2:0 format which implies that the chroma channels are being downsampled by a factor two in both the horizontal and in the vertical direction.

In the example shown, block **b0** contains 16 coefficients: one DC (the first one at coordinate 0,0) and 15 AC coefficients (the rest of the block). The DC represents the so-called “constant” value of luminance (for ‘Y’ blocks) and of the chrominance (for ‘U’ and ‘V’ blocks), and the ACs form the variable part meaning their contribution for each pixel is different. The combination of the DC and all ACs are used to represent the value of each pixel after decoding based on the used transform. The 16×16 luma frequency domain coefficients ‘Y’ are fully utilized to represent 16×16 spatial domain pixels. In the explanation above, the chrominance ‘Us’ are sub-sampled. This format is known as YUV 4:2:0, which means that four luminance pixels for each 2×2 square of the Y pixels share one ‘U’ and one ‘V’ pixel respectively.

There also exist other formats known as YUV 4:2:2 or YUV 4:4:4, in which the chrominance is not sub-sampled at all in YUV 4:4:4, or where the chrominance has twice as much data in 4:2:2. The present invention can work for any of these formats.

In the described example, the blocks **b0** to **b23** for the macroblock are sorted based on a measure (index value) of block energy or activity. FIG. 3a illustrates an example of block-sorting. The block energy used to order the sort can be measured in a number of different ways. According to one technique, the sort is based on the number of zero value coefficients in a block. In another technique, the sort is carried out using the average value of the modulus of non zero coefficients. Using a measure Δ of block energy, a search range is established within the sorted list as illustrated in FIG. 3 to identify candidate blocks (Step S4 of FIG. 2). The best candidate for prediction is then established as described above based on bit rate evaluation (Step S6 of FIG. 2).

The right hand diagram in FIG. 5A illustrates the effect of these predictions. Block **b12** is labelled **P1** to denote it as the first predicted block. Instead of transmitting the actual coefficients in block **b12**, coefficients (residuals) representing the differential between block **b12** and **b10** are transmitted, together with the information that in the transmitted data block **12** has been predicted from reference block **10**. An indication of the reference block **10** is also transmitted, e.g. identified by its frame number and position in the frame. This is shown schematically in the list on the right hand side of FIG. 5A where **P1** denotes prediction **1**, block **12** minus block **b10** in the luma block. The next candidate to be selected is block **20** labelled **P2** which is predicted from block **b21**. The process continues and in this case results in 7 predicted blocks. This results in a reduction in the number of coefficients to be transmitted by 9 (from 132 to 123). In a specific example, when the video data is encoded for transmission in bins, this has the effect that bins **122** and **1008** are removed,

16

while the content of bins **224** and **288** are increased. In FIG. 5A, the arrows denote the relationship between a predicted block and the block from which it is being predicted.

FIG. 5B shows a prediction example for motion prediction between different blocks of different macroblocks of two frames.

FIG. 6 is a schematic block diagram of a decoder for decoding a video stream which has been subject to the block prediction technique described above. In addition to the encoded coefficients, the video stream includes data defining the predicted blocks, the identity of the blocks from which they have been predicted and the order in which they have been predicted. The encoded stream is supplied to an entropy decoder **12** which determines for the incoming data whether the blocks to be decoded are for reconstruction of an intra frame or reconstruction of an inter frame. Blocks for reconstruction of an intra frame are passed to intra reconstruction stage **14**, while blocks intended for reconstruction of an inter frame are passed to inter reconstruction stage **16**. A predicted block is reconstructed by adding the residuals to the correspondingly located coefficients in the block it is predicted from. The output of the reconstruction stages **14** and **16** are supplied to an inverse quantization stage **18** and then to an inverse transform stage **20** where the quantization coefficients are transformed from the frequency domain into the time domain as a decoded stream.

Details of a preferred technique for matching blocks based on bitrate contribution are now discussed in more detail. This technique decreases the bitrate in video compression by means of block prediction in the quantized domain. The input to the method is e.g. a slice or a set of slices of blocks of transformed and quantized coefficients (e.g. residuals from the H.264). A slice means a group of macroblocks, so one slice per frame means all macroblocks in the frame belong to the slice. For each transformed and quantized block in the current slice a block from previous encoded slices or a block in the current slice (care has then to be taken to ensure a decodable stream) is a potential candidate to be used for prediction in order to reduce the bitrate (compared to direct entropy coding of the block itself). An example embodiment of a predictor and the “optimal” selection of the block to be used for prediction and required side-information to identify that block (needed description for reconstruction in the decoder) is described below. The side information is entropy encoded into the encoded bitstream along with the residual, by entropy encoder **10**.

In the preferred embodiments, the present invention performs block matching using two classes of metrics: one based on an aggregate or pattern of the block (e.g. energy, structure etc.) and a second based on bit rate. These two metrics are used in two separate stages: the first stage to sort and the second stage for the RD loop. In particularly preferred embodiments, the RD loop rate target is not only to find two blocks that can predict each other closely in terms of rate, but also to solve this problem for groups of blocks at the same time. One simple example could be the following patterns—(a) 1,2,1,2,1,2,1,2 and (b) 46,47, 46,47, 46,47, that will result in (a) 1,2,1,2,1,2,1,2 and (b) 45, 45, 45, 45, 45, 45, 45. That is to say, multiple blocks can be matched from the same sorted list of candidate blocks, including potentially both interframe and intraframe prediction being performed based on the same sorted list.

The advantages are improved entropy coding due to improved redundancy removal prior to an arithmetic or Huffman coder in the entropy encoder **10**. Compared to VC-1 [1, pp. 251] there are a number of potential advantages: (i) all coefficients in the block are used in the prediction and not just

the first row and/or column; (ii) all blocks in the frame/slice as candidate blocks for prediction and not just block to the left or on the top; (iii) generalized prediction structures, e.g. weighted prediction or prediction from multiple blocks; and (iv) explicit rate estimation for finding the best block for prediction (taking the cost for side information into account).

Let $X(m, n)$ denote a block mEM (a frame/slice consists of M blocks in total) of quantized coefficients (e.g. quantized DCT coefficients) at time-instance n . The blocks are conventionally fed to an entropy coder 10 (in H.264 more specifically the context adaptive variable length coder or the context adaptive binary arithmetic coder). That is, from the point where we have $X(m, n)$, lossless compression is performed, i.e., the distortion is fixed. The method seeks to remove remaining redundancies (and thereby reduce the rate) prior to the arithmetic coder by means of a predictor. In one embodiment the prediction is formed as a subtraction between a current block and a reference block. The optimal indices (o_{opt} , p_{opt}) for prediction of the current block $X(m, n)$ is selected based on rate calculation/estimation, i.e.,

$$(o_{opt}, p_{opt}) = \arg \min_{o, p} (R(X(m, n) - X(o, p)) + R(o, p)) \quad [1]$$

where $R(X(m, n) - X(o, p))$ denotes the bitrate of the prediction residual and $R(o, p)$ the bitrate of side-information (i.e., the bitrate for transmission of the prediction block index o of frame p). The rate estimation can for instance be provided from parts of the arithmetic coding routine where the sum of \log_2 of the symbol probabilities can be used to estimate the rate. It could also be beneficial, from e.g. a computational aspect, to approximate the criterion in equation [1] by using another measure that correlates well with the rate. Generally, any metric can be used that relates in some way to a number of bits that would be required in the encoded bitstream to encode both the residual block and the side information identifying the respective reference block (i.e. would be required for each candidate if that candidate was chosen as the reference block), whether the metric is a direct measure the number or rate of bits or a metric that correlates with the number/rate.

The search for the optimal predictor can be made computationally more efficient by pre-ordering the candidates such that potential good candidates are located in the proximity of a specific position in an ordered array. Let $Y(k, n)$ now denote element k in an M dimensional ordered array of the block indices mEM of frame n according to some measure. For instance, the blocks $X(m, n)$, $m \in M$ can be ordered according to their ascending energy (or some other signal dependent properties).

To find the predictors we start e.g. with the first block in the current frame/slice in the raster-scan order (or some other order beneficial from either a rate or computational complexity point of view) and find its position in the ordered array $Y(n)$ of the current frame and the (re-)ordered arrays of the previously processed frames $Y(n - \text{NumRef}), \dots, Y(n - 1)$. NumRef is the number of reference frames, i.e. here the number of previous quantized frames that have been processed and can be used for inter prediction. As prediction candidates from the current frame/slice we select the candidates that are located within the range $\pm W$ around the current index in the ordered array, i.e., the "intra" prediction candidates plugged into expression (1) are the blocks corresponding to the sorted indices $Y(q(n) - W, n), \dots, Y(q(n) - 1, n)$; and $Y(q(n) + 1, n), \dots, Y(q(n) + W, n)$; where q denotes the position of the current block in the ordered array. Note that special caution has to be taken to avoid cyclic predictions, i.e., avoid prediction of block m from block n if block n has already been predicted from block m , making decoding infea-

sible. It should also be mentioned that direct encoding (i.e., no prediction) of the residual is also included as a candidate for the rate estimation.

Similar to the selection of candidates for intra prediction, the inter prediction candidates are selected as $Y(q(n-i) - W, n-i), \dots, Y(q(n-i) + W, n-i)_{i=1: \text{NumRef}}$.

All intra and inter candidates are evaluated according to equation (1) and the optimal index pair is selected. This procedure is repeated for all blocks in the frame/slice. The resulting prediction residuals (variable/index differences) together with required side-information for decoding is e.g. arithmetically encoded and sent to decoder.

Referring to FIG. 10 one embodiment of the method performed by the encoder is as follows.

Step T1: order all the blocks in the frame according to some measure.

Step T2: set block index to $m=0$.

Step T3: find the equivalent position q of the block index m in the ordered lists (both current and previous quantized frames, i.e., find $q(n), \dots, q(n - \text{NumRef})$).

Step T4: select the intra and inter prediction candidates as $Y(q(n) - W, n), \dots, Y(q(n) - 1, n)$; $Y(q(n) + 1, n), \dots, Y(q(n) + W, n)$; and $Y(q(n-i) - W, n-i), \dots, Y(q(n-i) + W, n-i)_{i=1: \text{NumRef}}$ respectively.

The size of the search range W is a trade-off between performance and computational complexity.

Step T5: find the best candidate according to expression [1] or some approximation of it.

Step T6: send optimal prediction residual together with side-information (e.g. the position of the residual block within the frame and the position (e.g. space and time) of the block that was used for prediction) to the arithmetic coder.

Step T7: increment block index $m=m+1$ and go to step T3, until $m=M$ when the method moves to the next frame $n=n+1$.

Referring to FIG. 11, one embodiment of the method performed by the decoder is as follows.

Step U1: decode all prediction residuals and side information (this gives a frame of prediction residuals together with the description for each block how to undo the prediction).

Step U2: reconstruct all blocks that do not depend on unreconstructed blocks (i.e. (undo prediction)).

Step U3: repeat step U2 until all blocks have been reconstructed.

The above example embodiment can be extended in several ways. For instance it could be beneficial to use weighted prediction or prediction from multiple blocks. Additional side information would be needed to be transmitted which for weighted prediction and prediction using multiple blocks would be prediction weights and block positions/addresses.

An illustration of the prediction in the encoder is shown in FIG. 9. This gives a high-level illustration of the block prediction in the encoder. The prediction residual together with side information is sent to the entropy coder. In the decoder the reverse procedure is performed, i.e. first reconstruct the residual frame and then reconstruct the frame given side information.

The above described embodiments of the present invention may provide several advantages. Matching blocks are located by examining the difference between blocks to be certain that the bit rate of the ultimately transmitted video data will be reduced with respect to the bit rate for sending coefficients of those blocks. Moreover, the pre sort has identified candidate blocks within which this comparison takes place. The blocks do not have to be physical neighbours in the image frame—instead, they are sorted on the basis of an index value associated with the blocks, for example, representing energy. This

allows a best matching block to be selected from any part of a frame (or indeed a different frame). When selecting best candidates, the comparison of bit rates can take into account the overhead information that needs to be transmitted to identify that the block is a predicted block, and to identify the block from which it is predicted. The identity of the block from which it is predicted can be provided to the decoder in the form of an location within the frame expressed as a number of intervals of whole blocks, i.e. a block address, rather than by a motion vector expressed as an offset in terms of a number of pixels or even fractions of pixels.

The method described removes redundancy in the temporal and frequency domain before and/or after quantization in a compressed digital video stream by means of block prediction. The input to the method is a set of transformed and/or quantized transform coefficients of a set of frames in the temporal domain of the input video stream. The input video stream frame can be separated into blocks and groups of blocks. The groups of blocks are not limited by the location of the individual blocks participating in the group. The prediction is performed between the blocks of the current frame (intra) and is not limited by location of the blocks but by the factor of the block similarity. The same technique can be used for inter frame predictions. Inter frame block matching is not restricted by location either. The block similarity is determined from the point of view of reduction of bit rate.

Furthermore, as explained, in a preferred embodiment processing is carried out in the frequency domain where the transform has already compacted the energy of the target object such that comparison can now be carried out using a few frequency domain coefficients instead of a whole image. In these embodiments, both components of the method, i.e. processing in the frequency domain and the sort versus search, reduce the complexity while maintaining a very high quality. Another benefit of the method is that due to the low complexity of the calculations involved, the number of used reference frames for inter frame motion compensation can be substantially higher in comparison with existing algorithms. Another major benefit is that, due to the low complexity, matches can be made on several level sub block divisions. That is, an image portion can be a macroblock, a block or even a smaller number of pixels than a block. This is because the described method achieves low complexity and therefore incurs fewer of clock cycles, which if desired means that some of the saved complexity can then be spent searching for sub-blocks such as 4x4 or 2x2 sub-blocks instead of just blocks. Alternatively the search could be performed at a higher level of 16x16, 32x32 or 64x64 aggregate blocks for example, which would save on the side information necessary to signal them in the encoded stream.

A particular advantage arises from processing in the frequency domain. While there are frequency domain processing models existing, there is none that explores redundancy reduction as in the method described in the following embodiments; in particular which provides a unique set of benefits including complexity reduction, preserving and improving video quality and at the same time lowering the bit rate of the encoded video stream.

The method need not require loop filter or loop back to the spatial domain for motion estimation due to the fact that all processing is now concentrated in the frequency domain. This is a major advantage with respect to existing video coding methods and a point of significant reduction of complexity.

Another advantage is that processing of all the colour components can be done at the same time. That is, processing

done in the luminance channel can affect processing done in the chrominance channels.

Another advantage of processing in the frequency domain relates to blocks lying on the edge of a frame or slice of a sub frame. That is, the blocks that lie on the edge of a frame (or if a sub frame separation in multiple slices is used, the blocks that are on the edge of the slice) can be efficiently predicted. As the blocks are grouped in accordance with similarity, the method allows grouping of blocks or slices in any order and hence there is no penalty in the prediction of blocks sitting on the edge of a slice or frame. This is a significant improvement in comparison with the current FMO (Flexible Macroblock Ordering) in the current Standards like MPEG-4 AVC/H.264.

Another advantage of the described embodiments of the invention herein is that deep sub-block sub-divisions can be utilised without excessive processor load.

Note that the different preferred techniques discussed above need not necessarily be used in conjunction with one another. For example, it is possible to perform block matching in the frequency domain without using the additional technique of a sorted list based on block energy or other such index. Alternative block matching techniques could also be used, for either intra and/or inter frame block matching, e.g. by matching based on a measure of correlation or a measure of minimum error. Conversely, it is possible to use the sorting technique for block matching without a frequency domain transform, e.g. by determining a measure of block energy based on the spatial domain coefficients (though this is less preferred since it will tend to be more computationally intense).

Further, where sorting is discussed as a method of determining a subset of candidates within a search range Δ , note that it is not necessarily required to rearrange list entries in memory. More generally, the search for candidates may be performed by any method of identifying blocks having an energy or other index within the desired range.

Further, the sort index need not necessarily be a measure of block energy. Another possibility would be a metric relating to the structure of the block, such as the structural similarity index (SSIM). In other embodiments, multiple metrics could be combined in order to determine the index used for sorting. Furthermore, once the list is sorted, aspects of the invention need not necessarily be limited to finding the best match from amongst the candidates based on bitrate contribution. Other second metrics could be used for this purpose, e.g. a more conventional motion based matching as used in H.264.

Signalling Blocks by Address

The above describes a particularly advantageous method of selecting reference blocks from a frame; but regardless of how blocks are selected from frames, the present invention provides an improved method of encoding the identity of reference blocks for transmission to the decoder. Exemplary details of a method of signalling the intra and inter prediction information for prediction in the frequency domain are now discussed in more detail.

As mentioned, according to one aspect of the present invention a block is matched only to another whole block rather than to a block-sized area offset by any number of pixels as in more conventional block matching techniques. Therefore the signalling algorithm of the present invention sends block addresses instead of motion vectors, i.e. represented in terms of a whole number of blocks rather than a pixel offset. Note however that whilst the term "block" may be used herein, in its most general sense this is not intended to

21

imply and particular size, shape or level subdivision. It will be appreciated that in different schemes then various different divisions and subdivisions may be referred to by terms such as macroblock, block and sub-block, etc., but that the term “block” as used most generally herein may correspond to any of these or indeed any other constituent image portion being a division of a video frame corresponding to multiple pixels. Whatever manner of division is employed, according to the present invention the address of the reference portion for use in prediction is signalled as a whole number of multi-pixel portions instead of a pixel offset.

In embodiments, the bitstream may also contain one or more prediction method flags indicating a prediction method to be used by the decoder (corresponding to that used by the encoder).

Further, the bitstream may contain a frame number of the reference block, as the reference block for prediction can be chosen from any of multiple different frames.

In one particularly preferred embodiment, the side information signalled in the bitstream to the decoder comprises:

22

frame number, an addition or subtraction flag, absolute value flag, a macroblock address, a block address within the macroblock, and a sub-block address within the block. The signalling structure of this side information is shown in the following table.

Field	No. Bits
Frame Index (FrameIdx)	4
Add/Sub	1
Nat/Abs	1
Macroblock Address (MBAAddr)	9
Block Address (BlockAdr)	3
Sub-block Address (SubBAdr)	2

For each resolution the prediction indexes can be encoded as follows. This shows signalling structure size and encoding for a plurality of different resolutions.

Inter	SIF	WVGA	480p	4CIF	720p	1080p	4k × 2k	8k × 4k	Values	Max	4 bits
FrameIdx	4	4	4	4	4	4	4	4	−8 . . . 7	0	Intra
Sub/Add	1	1	1	1	1	1	1	1	0 . . . 1	1	−1
Nat/Abs	1	1	1	1	1	1	1	1	0 . . . 1	2	−2
MBAAddrX	5	6	6	6	7	7	8	9	0 . . . Max	480	3 −3
MBAAddrY	4	5	5	6	6	7	8	9	0 . . . Max	270	4 −4
BlockAdr	3	3	3	3	3	3	3	3	0 . . . 5	5	−5
SubBAdr	2	2	2	2	2	2	2	2	0 . . . 3	6	−6
Total/B	20	22	22	23	24	25	27	29		7	−7
Total/MB	120	132	132	138	144	150	162	174		−8	List

This improved prediction scheme is more effective than the current prediction scheme which use a higher bit rate to signal only part of the information that the improved scheme can transmit. The streamlined inter and intra prediction allows for simplified signalling method. FIG. 3b shows a block matching prediction example achieving bit savings. The table below shows the effective side information and coding for multiple resolutions.

	Res X	Res Y	MB_x	MB_y	MBs	MBBits	UpToMBs	Bs	BBits	UpToBs	Bits_X	Bits_Y	Bits_XY
SIF	320	240	20	15	300	9	512	1800	11	2048	5	4	9
WVGA	640	400	40	25	1000	10	1024	6000	13	8192	6	5	11
480p	640	480	40	30	1200	11	2048	7200	13	8192	6	5	11
4CIF	704	576	44	36	1584	11	2048	9504	14	16384	6	6	12
720p	1280	720	80	45	3600	12	4096	21600	15	32768	7	6	13
1080p	1920	1080	120	68	8160	13	8192	48960	16	65536	7	7	14
4k × 2k	3840	2160	240	135	32400	15	32768	194400	18	262144	8	8	16
8k × 4k	7680	4320	480	270	129600	17	131072	777600	20	1048576	9	9	18

Signalling Blocks by Means of a “Global Block List”

The following describes a second improved method of encoding the identity of reference blocks for transmission to the decoder. Again, this may be used regardless of how blocks are selected from frames. In embodiments this method may furthermore extend the available candidates to include certain “notional” or “artificial” blocks (rather than just actual blocks appearing in actual frames).

As discussed with reference to FIG. 12, according to this aspect of the present invention the encoder generates a table of prediction blocks (i.e. reference blocks for use in prediction) having the most regularly (often) used block characteristics. A block in the table can then be referenced in the encoded signal instead of a block in the frame. The table of most common blocks is determined during encoding and will be updated dynamically for transmission to the decoder.

Thus the encoder generates an ad-hoc codebook for signalling the reference block to the decoder.

For example, certain blocks such as those shown in FIG. 1b, 5A or 5B may occur more regularly within a certain frame, sequence of frames or part of a frame to be encoded (i.e. a higher number of instances). If certain blocks (i.e. certain sets of block coefficients or approximate sets) occur often enough, then it may become more efficient to dynamically maintain and transmit to the decoder a look-up table of these regularly encountered blocks and then signal the identity of reference blocks used in the prediction coding by reference to an entry in the look-up table, rather than identifying the block by some other means such as a motion vector or block location address.

Table entry #	Block definition
0	$B_a(\dots)$
1	$B_b(\dots)$
2	$B_c(\dots)$
3	$B_d(\dots)$
4	$B_e(\dots)$
etc . . .	

Each block definition indicates a certain respective regularly-encountered set of block coefficients (or approximate set).

Preferably the table is updated and transmitted periodically, and each updated table may be interleaved into the encoded data stream as shown in FIG. 12 (though other methods of separate transmission are not excluded). Thus after a certain table L_n is transmitted in the bitstream to the decoder, then one or more subsequent encoded blocks (n,1), (n,2), etc. that are each encoded by prediction coding based on another reference block are each transmitted in conjunction with side information $S_{n,1}$, $S_{n,2}$ indicating a respective entry in the look-up table. When the look-up table is updated and retransmitted (L_{n+1}), then subsequent side information $S_{n+1,1}$ may then signal a reference block for use in prediction by indicating an entry in that updated table. The decoder stores a copy of the most recent look-up table L received in the bitstream and uses it in conjunction with the side information S to identify a reference block for use in predicting a current block to be decoded (and combines the predicted block with the respective residual data).

This technique is particularly useful when block matching is performed in the frequency domain, because the energy of each block compacted into only a few non-zero coefficients—for example see FIG. 1b. In this case certain blocks are likely to be selected in the block matching process often enough for the maintenance and transmission of a dynamic

look-up table to be an efficient choice of technique. Nonetheless, the described technique is not limited to frequency domain processing nor selection based on block sorting, and could also be used to encode the results of other block matching techniques.

Note that in most practical embodiments the look-up table will not be exhaustive. I.e. some reference blocks will not be selected very regularly by the block matching process, and those blocks will not be included in the look-up table. Such reference blocks may be signalled in the encoded bitstream by a different method, e.g. preferably in the form of a block address identifying an absolute block position in the frame or a relative block address between the current block and respective reference block. That is, a block location expressed in terms of a whole number of multi-pixel blocks (rather than any number of pixels or fractional number of pixels as in a conventional motion vector).

In embodiments the actual coefficients of the most regular reference blocks may be signalled to the decoder in the look-up table L (though an alternative would be for the encoder and decoder to have all possible block definitions that could potentially be signalled in the table pre-stored at either end).

Thus according to the above techniques, a subset of blocks frequently appearing in multiple frames of matched blocks will be separately transmitted to the decoder in a group. In a particularly preferred embodiments, this group may not only include existing blocks but can also include “artificial” blocks with sets of coefficients that help the prediction process and are calculated from the blocks in the input stream. That is, one some or all of the blocks in the group need not be actual blocks of the frame, but instead could be notional blocks comprising pre-defined “artificial” coefficients which may be set by the system designer or calculated by the encoder, e.g. by averaging, interpolating or extrapolating from other actual blocks regularly found in the video. In the preferred embodiments these artificial blocks would be included in the sorted list or as candidates in the block matching process.

For example, say a number of blocks regularly occur having a particular frequency domain coefficient which is regularly within a certain range, e.g. 200 to 210. In this case the encoder may create an artificial block having a corresponding coefficient with an average or interpolated magnitude within that range, e.g. 205. That way, the regular similar blocks encoded relative to that artificial block will result in only a small residual, e.g. typically no more than about 5 in size, thus reducing the bitrate required to encode that residual. This example of a particular coefficient is considered for the sake of illustration, but note that the greatest benefit of this scheme will be achieved when blocks regularly occur with similar patterns of multiple coefficients, in which case an artificial reference block can be generated having an approximate pattern close to that of a number of those blocks (the single coefficient case is likely to be handled well in the entropy encoder anyway). It is these multi-coefficient blocks which tend to incur most of the bitrate in the encoded signal.

The group of blocks which populate the look-up table may be referred to herein as the “global group”, which may comprise artificial blocks and/or actual blocks which are extracted from one or more actual frames of the video. As described, the global group is updated dynamically based regularity of use in the prediction coding.

The best candidate is preferably selected as the lowest bit rate contributor, selected from the set of blocks in the global group and the actual blocks in the frames in any combination.

25

One unique aspect of this algorithm is in the approach of application of the global group of blocks to aid the prediction process. The extraction of existing blocks to the global group will further reduce the bitrate.

The members of this group will be defined during the encoding process and will be based on rate calculation (i.e. how many bits would be required in the encoded bitstream), thus creating a set of coefficients that are the best predictors for a set of blocks in one or multiple slices of frames and not necessarily existing blocks.

Due to the rate prediction nature of the process then a perfect match is not necessarily sought, but rather rate reduction which allows for finding the closest match instead of the exact pattern.

All of the blocks in the frame or only the blocks submitted as the global list can be shifted in this way, creating another opportunity for bitrate reduction. Also, slices of the frame can be shifted as well as parts of the global group.

The look-up table will be periodically updated by allowing the decoder to drop blocks that will no longer be required from the table, and update them with new ones from the incoming stream.

In practice, the look-up table may not always be a useful way of signaling reference blocks, e.g. if the encoder determines that few or no blocks are selected significantly more than any others. In this case, the look-up table may not actually save on bits in the bitstream since it will be rarely referenced whilst itself incurring a number of bits in overhead. Therefore in particularly advantageous embodiments, the encoder may be configured to detect some bit-saving threshold, and if the saving is small or even negative then it will cease using the look-up table method (and stop sending updates), and instead will signal reference blocks by another means such as identifying the location of the reference blocks by their address within the stream. In order to inform the decoder which encoding method is used, the encoder will also include a flag in the bitstream indicating whether or not the look-up table method is currently being used. The flag may also be sent periodically, or on an ad-hoc basis as and when the encoder decides to change mode.

The algorithm described here is lossless. The encoding and the decoding process can be made very low complexity algorithms since the process operates in the frequency domain on a very small set of coefficients (in comparison with the spatial domain process where the complexity is exponentially higher).

Scaling and Rotation

One flaw with conventional codecs which perform motion estimation in the spatial domain is that they require high computational complexity in order to handle prediction based on non lateral (non translational) motion, i.e. scaling or rotation. Scaling occurs when the camera zooms in or out, or an object moves closer or further away from the camera, or indeed if the object expands or shrinks. Rotation occurs when the camera and object or background rotate relative to one another.

An advantage of performing motion estimation in the frequency domain is that the complexity of handling scaling and rotation type prediction is greatly reduced.

For example, consider the illustrative example of FIG. 13a. On the left hand side is represented a block B at some point in time, comprising a frequency domain coefficient C. That is, the block may be considered to comprise a frequency domain term such as a sinusoid with amplitude C which varies with some wavelength across the block, e.g. representing a variation in chrominance or luminance across the block. An example of a sinusoidal variation in chrominance or lumi-

26

nance is illustrated in the top left of FIG. 13a and the corresponding array of coefficients used to represent the block is shown in the bottom left. Of course other frequency domain terms may also be present, but for illustrative purposes only one is shown here.

Now consider a corresponding block B' at a later point in time when the camera has zoomed out from the object or scene in question (or the object has moved further away or shrunk). As shown in the top right of FIG. 13a, this means the wavelength of the frequency domain term decreases. Suppose for the sake of an illustrative example that the zoom causes the wavelength to halve: the effect in the frequency domain representation of the block is that the coefficient C moves from one position to another, as shown in the bottom right of FIG. 13b. That is, the energy of the block is redistributed from one frequency domain coefficient of the block to another. In reality there is unlikely to be a sudden zooming out to exactly half the wavelength, but the zoom may still result in a gradual "fading" transition of the block energy from some coefficients to others—i.e. coefficients representing the amplitude of lower frequency terms will gradually decrease while coefficients representing higher frequency terms gradually increase. E.g. the coefficient at the second (right-hand) position in FIG. 13a will gradually increase at the expense of the coefficient at the first (left-hand) position. A similar effect will occur in the opposite direction for a zooming in (or the object moving closer or expanding).

Given such a scenario, the scaling motion of one block can be predicted from another block and may be encoded relative to that block using only very few bits. It is the transformed representation of the blocks that allows this prediction and encoding to be achieved with a low bitrate and low computational burden, because in the frequency domain the scaling will typically involve only a very low-complexity transition of frequency domain terms.

In one scenario the scaling may just be encoded in terms of the difference between the frequency domain coefficients of one block and another, on the basis that as the energy of the block gradually fades from one frequency domain coefficient to another then the residual between one block to the next will be small. According to a further aspect of the present invention however, the encoder may estimate some parameter of the scaling and signal that parameter to the decoder as side information in the encoded bitstream. For example the parameter may comprise an indication of a scaling factor indicating that the target block can be predicted by scaling the selected reference block by some factor $\pm S$ (e.g. a certain percentage zoom in or out). The scaling may be estimated by a suitable image analysis algorithm in the encoder or by an auto-focus feature of the camera.

In the case where an estimation of the scaling factor is used in the encoding, the residual will represent only the difference between the scaled prediction and the actual target block, and will be even smaller so require even fewer bits to encode. That is, the encoder estimates a local prediction of the scaling to be applied to the selected reference block, subtracts the frequency domain coefficients of the scaled reference block from those of the target block so as to generate a residual (in practice this may just involve comparing the coefficients from shifted positions within the block), and then encodes the target block in the form of the residual, an indication of the scaling parameter, and an indication of the reference block. The signalled scaling parameter enables the decoder to determine the shift or transition to apply to the frequency domain coefficients of the reference block in order to recreate the predicted version. Adding the frequency domain residual at

the decoder then recreates the target block. Alternatively, in a non-lossless case the residual may be omitted from the encoding and decoding altogether.

Similar lox-complexity bitrate savings can be achieved by predicting rotation in the frequency domain. An illustrative example is shown in FIG. 13b

In this example, the left hand side again represents a block B at some point in time, having a frequency domain coefficient C representing a frequency domain term such as a sinusoid with amplitude C which varies with some wavelength in the horizontal or vertical direction across the block (e.g. representing a variation in chrominance or luminance). An example of a sinusoidal variation in chrominance or luminance is illustrated in the top left of FIG. 13b and the corresponding array of coefficients used to represent the block is shown in the bottom left. In this case, the coefficient C in the first block B represents a certain variation in the horizontal direction across the block. Again note that other frequency domain terms may also be present, but for illustrative purposes only one is shown here.

Now consider a corresponding block B' at a later point in time when the camera or object has rotated by 90°. As shown in the top right of FIG. 13b, this means the frequency domain term is flipped from a horizontal to a vertical orientation. The effect in the frequency domain representation of the block is that the coefficient C is flipped about the diagonal axis from one position to another, as shown in the bottom right of FIG. 13b. That is, the energy of the block is redistributed from one frequency domain coefficient of the block to another corresponding coefficient representing the same frequency term but in the transverse direction across the block.

In reality there is unlikely to be a sudden right-angled rotation. However, the effect can be generalised to other angles of rotation. This takes advantage of the fact that the block or macroblock tends to have approximately the same total or average energy when rotated—i.e. the object in question is not emitting more light, just changing orientation relative to the camera. As shown in FIG. 13c, in a generalised rotation, the target block (centre) of some later frame F' can be predicted based on a reference block from an earlier frame F, and the contribution to the block energy from the neighbours may be approximated to be small (the contributions being shown shaded black in FIG. 13c) and/or similar to the energy lost to the neighbouring from other regions.

Hence as the image rotates, the energy of the block will gradually fade between coefficients. That is, the energy from the coefficient at one position is gradually redistributed to its diagonal counterpart. E.g. in the following rotation the factors “a” and “b” are given by the calculation illustrated in FIG. 13e.

$$\begin{array}{ccccccc} \dots & \dots & C & \dots & \dots & \dots & b^*C \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \rightarrow & a^*C & \dots \dots \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \end{array}$$

That is, $a=\sin(\alpha)$ and $b=\cos(\alpha)$ where α is the angle of rotation. Preferably, the encoder will estimate the rotation using one of a set of computationally relatively low complexity rotations such as 30°, 45°, 60° and 90° as a best approximation.

For 30°, $a=1/2$, $b=(\sqrt{3})/2$

For 45°, $a=1/(\sqrt{2})$ and $b=1/(\sqrt{2})$

For 60°, $a=(\sqrt{3})/2$, $b=1/2$

The residual may then encode the difference between the approximated predicted rotation and the actual block.

In one scenario the rotation may just be encoded in terms of the difference between the frequency domain coefficients of one block and another, on the basis that as the energy of the block gradually fades from one frequency domain coefficient to another then the residual between one block to the next will be small. In a further aspect of the present invention however, the encoder may estimate some parameter of the rotation and signal that parameter to the decoder as side information in the encoded bitstream. For example the parameter may comprise a rotation angle indicating that the target block can be predicted by rotating the selected reference block by the specified angle. The rotation may be determined by an image analysis algorithm in the encoder or by gyro sensors of a mobile terminal in which the camera is housed.

In the case where an estimation of the rotation angle is used in the encoding, the residual will represent only the difference between the rotated prediction and the actual target block, and will be even smaller so require even fewer bits to encode. That is, the encoder generates a local prediction of the rotation to be applied to the selected reference block, subtracts the frequency domain coefficients of the rotated reference block from those of the target block so as to generate a residual (in practice this may just involve comparing rows of coefficients with columns and vice versa), and then encodes the target block in the form of the residual, an indication of the rotation angle, and an indication of the reference block. The signalled rotation parameter enables the decoder to determine the flip or transition to apply to the frequency domain coefficients of the reference block in order to recreate the predicted version. Adding the frequency domain residual at the decoder then recreates the target block. Alternatively, in a non-lossless case the residual may be omitted from the encoding and decoding altogether.

In embodiments the encoder will have the option of using any of the lateral, scaling and rotational types of motion prediction for encoding any given target block. In that case, it is useful to provide a mechanism for selecting the type of prediction to use for each target block. One such mechanism is, for a group of potential reference blocks, for the encoder to try each type of prediction in turn according to a type-hierarchy. Preferably, the encoder first attempts a lateral (i.e. translational) type prediction based on each of the candidates (this typically being the least computationally complex type of motion prediction). If a match is found which will result in a bitrate contribution within a maximum threshold, then the lateral type of prediction used based on that match, and the matching process halts there so that the scaling and rotation type predictions are not even considered for that target block. That is, if the number of bits required to encode the residual plus side information for the target block based on the best matching reference block is found to be within a certain threshold using a lateral type prediction, then the lateral prediction is used and other types are ignored so as to try to avoid wasting unnecessary machine cycles. However, if no match is found which would provide a bitrate contribution within the threshold using lateral type prediction, then one of a scaling or rotation type prediction may be tried. E.g. the next in the hierarchy may be scaling. The encoder therefore attempts a scaling type prediction based on each of the candidates in the list and tests whether the best matching candidate falls within the maximum bitrate contribution threshold if scaling is used instead of lateral prediction. If so, it encodes the target block based on the best matching reference block using scaling type motion prediction and halts the block matching process for that target block. If not however, the encoder then attempts a

rotational type prediction for each candidate in the list and tests whether the best match falls within the bitrate contribution threshold using rotation type prediction. If so, it encodes the target block accordingly. If no matching candidate is found within the maximum bitrate contribution threshold, the encoder may accept the best of a bad lot, or may extend the list of candidates, or may encode by conventional intra encoding or encoding of absolute values.

In other embodiments the hierarchy may be different, e.g. lateral, rotation, scaling. An alternative would be to compare all types of prediction together without hierarchy. However, that would incur a high processing burden and is less likely to be desirable for a live video stream.

The encoder may signal an indication of the type of prediction in the encoded bitstream, so the decoder knows what type of prediction to apply. As mentioned, the encoder may also signal a parameter of the motion. In the case of rotation, the signalled information may indicate a degree of the rotation. In the case of scaling, the signalled information may indicate a scaling factor. This information allows the decoder to reconstruct a prediction of the target block based on the signalled rotation

Referring to FIG. 13d, prediction of rotation in the frequency domain can be particularly advantageous when combined with the feature of selecting reference blocks from a sorted list, as discussed above. FIG. 13d shows a screen or video window 50. As shown, if a large area is rotated then a closely matching candidate B for rotation type prediction of a target block B' may in fact be found a large distance away within the screen or viewing window 50. A conventional codec which predicts blocks only based on spatially neighbouring regions of the image would miss this situation. However, using the sorted list according to certain aspects of the present invention, blocks from any part of the screen may become candidates for prediction. Particularly when the list is sorted according to block energy, then blocks which closely resemble rotations of one another will become very close in the sorted list (regardless of distance from one another) since the rotation typically involves little variation in total block energy. Therefore a sorted list in which candidates are identified based on similarity of block energy is particularly likely to find good candidates for rotation type prediction.

The scaling and rotation types of prediction can also be particularly advantageously implemented using the feature of signalling reference blocks according to a global block list of the kind discussed above. In such an implementation, the reference blocks in the global block list can include representations of scaled and rotated patterns. For example, an artificial reference block may be generated which is suitable for prediction coding of multiple target blocks according to a number of different types of prediction.

E.g. consider an artificial reference block having energy condensed into the following two non-zero coefficients:

```
224 0 0 0
0 0 0 0
280 0 0 0
0 0 0 0
```

This can be used to encode a 4×4 target block according to any of the following predictions with a reduced bitrate.

90° rotation

```
~~~~~
224 0 0 0 224 0 280 0
0 0 0 0 0 0 0 0
280 0 0 0 → 0 0 0 0
0 0 0 0 0 0 0 0
```

45° rotation

```
~~~~~
224 0 0 0 224 0 140 0
0 0 0 0 0 0 0 0
280 0 0 0 → 140 0 0 0
0 0 0 0 0 0 0 0
```

Zoom out

```
~~~~~
224 0 0 0 224 0 0 0
0 0 0 0 0 0 0 0
280 0 0 0 → 0 0 0 0
0 0 0 0 280 0 0 0
```

An example of spatial correlation where scaling would be applicable is in the color representation of 4:2:0 and 4:2:2 formats where the scaling is defined by the color sampling.

This approach allows for prediction of scaled or rotated blocks due to the similarity of the pattern that each block is covering. The rotation or scaling will be expressed as a reordering of coefficients in the same block. Due to the rate prediction nature of the process then a perfect match is not necessarily sought, but rather rate reduction which allows for finding the closest match instead of the exact pattern.

Super Resolution

As shown in FIGS. 14a and 14b, it is possible when reconstructing an image at a receiver to overlay frames which are offset by a fraction of a pixel from one another in order to achieve a higher resolution. This idea may be referred to as "super-resolution".

FIG. 14a illustrates a pixel grid (raster) having some particular resolution defined by the pixel size of the camera. When a frame is captured, the image has a resolution of one pixel value per unit of the grid, i.e. one value per pixel (per statistic required to define a single pixel, e.g. one value of Y, one value of U and one value of V per pixel in YUV colour-space). Say now that the pixel grid is offset right by approximately half a pixel and down by approximately half a pixel, either because the camera moves slightly or because the scene or object being captured moves relative to the camera. If this movement is known or can be estimated, then it is possible to reconstruct a higher resolution image by superimposing the values captured from the two offset grids. In the example of FIG. 14a this results in four "super-resolution" pixels A, B, C and D for each actual physical pixel of the camera's sensor array. Each super-resolution pixel value may be determined for example by interpolating between the two overlapping real pixel values which contribute to it.

FIG. 14b illustrates the idea in more detail. For the sake of illustration, suppose there is a camera with only a 2×2 sensor array of four pixels, and in a first frame at some moment in time an object is captured appearing only in the top two pixels and not appearing in the bottom two. Thus in the first frame the object contributes only to the top two pixel values and not the bottom two. Suppose then that by a second frame at a later

moment in time the object has moved down by half a pixel, or the camera has moved up by half a pixel. When the object is now captured in the second frame, different areas of the object appear in all four pixels and so it contributes to all four pixel values.

The idea of superimposing fractional shifts in a pixel grid has been used in the past to increase the resolution of satellite images for example. As mentioned, this idea may be referred to as “super-resolution” or sometimes “remote sensing” in the context of satellite images. However, this technique has only been used in the past to increase the resolution beyond the intrinsic physical resolution of the camera or detector in question. For example, some satellite detectors only have one “pixel” with resolution of the order 1 km, and rely on this technique to greatly improve on the resolution of the satellite detector.

However, no-one has previously considered the potential to deliberately transmit a video image with a lower resolution than the intrinsic resolution of the camera, then use a super-resolution scheme to reconstruct an image at the receiver having a resolution more closely approaching the camera’s intrinsic resolution. It is this idea that is the subject of a further aspect of the present invention. The advantage is that the transmitted video stream requires fewer bits per unit time. That is, instead of using fractional shifts of real size pixels to boost resolution beyond the camera’s natural resolution, one aspect of the present invention instead uses the super-resolution technique to transmit averaged values for larger image units each corresponding to multiple real pixels (the averaged units thus having lower resolution than the real camera resolution) and to then reconstruct the real camera resolution at the receiver (or at least a higher resolution than that of the averaged units).

An example is discussed in relation to FIG. 14c. Here, a region of an image is captured having higher-resolution values A to P:

A, B, C, D
E, F, G, H
I, J, K, L
M, N, O, P

In some embodiments, these higher-resolution values may correspond to the values captured from individual pixels of the camera’s sensor array. In other embodiments however, these higher-resolution values need not necessarily correspond to the actual physical size of the camera’s pixels, but rather may represent the smallest size unit that would be used by the encoder in question in some particular mode of operation. The point is that the following encoder will encode a frame with an even lower resolution, i.e. by averaging or otherwise combining groups of higher-resolution values to create larger, lower-resolution units represented by respective lower-resolution values. In the example of FIG. 14c the lower-resolution units are 2x2 groups of higher-resolution values, but it will be appreciated that other schemes could equally well be used.

At an initial frame in a sequence, frame 0, the encoder averages the higher-resolution values F, G, J and K (or otherwise combines them, e.g. by totalling). This average provides a single overall lower-resolution value for a single, larger, lower-resolution unit covering the area of the respective group of four smaller, higher-resolution units. A similar averaging is performed for adjacent groups, thus generating a lower-resolution grid of larger size units represented by respective lower-resolution values. The encoder then encodes and transmits the frame based only on the lower-resolution grid of the averaged values.

Note that in embodiments, the image may still be divided into blocks and/or macroblocks, with each block or macroblock comprising a plurality of lower-resolution units (though fewer than if represented at the higher resolution). In this case, the blocks of multiple lower-resolution units may still be transformed into the frequency domain as part of the encoding process, though the transform may be considered optional according to this particular aspect of the invention. Either way, the super-resolution algorithm operates in the spatial domain (i.e. if there is a frequency domain transform, the super-resolution algorithm occurs before the transform at the encoder and after the reverse transform at the decoder)

At a first subsequent frame in the sequence, frame 1, the encoder shifts the lower-resolution grid up and left by one higher-resolution unit. The encoder then averages the higher-resolution values A, B, E and F to create a single overall lower-resolution value for a single, larger, lower-resolution unit covering the area of the respective group of four smaller, higher-resolution units—so now offset in each of the horizontal and vertical direction by one higher-resolution unit, which means a fractional offset of the lower-resolution grid. Again a similar averaging is performed for adjacent groups, thus generating a lower-resolution grid of larger size units each represented by respective lower-resolution values—but this time including the described offset. The encoder then encodes and transmits the frame using only the offset lower-resolution grid of the averaged values (again with a transformation of blocks of multiple such units into the frequency domain if appropriate to the embodiment in question).

Note that the receiver has now been provided with two lower-resolution units covering the higher resolution unit F by means of the fractional overlap between the lower-resolution units of frames 0 and 1, thus allowing the receiver to generate an individual higher-resolution value for F. However, further subsequent frames 2-5 will be required in order to recreate the full higher-resolution grid.

At a second subsequent frame, frame 2, the encoder shifts the lower-resolution grid up and right by one higher-resolution unit relative to that of the initial frame 0. The encoder then averages the group of higher-resolution values C, D, G and H to obtain a respective lower-resolution value for a respective lower-resolution unit, and similarly for surrounding units, thus generating another offset grid of lower-resolution values which is encoded and transmitted to the receiver. The receiver now has enough information to recreate higher-resolution unit G by means of the fractional overlap between the lower-resolution units of frames 0 and 2.

At a third subsequent frame in the sequence, frame 3, the encoder shifts the lower-resolution grid down and right by one higher-resolution unit relative to that of the initial frame 0, and then averages the group of higher-resolution values K, L, O and P to obtain a respective lower-resolution value for another lower-resolution unit. This is encoded and transmitted to the receiver as part of a grid of similarly offset lower-resolution units, now allowing the receiver to recreate higher-resolution unit K by means of the fractional overlap between the lower-resolution units of frames 0 and 3.

The sequence then continues to a fourth subsequent frame, frame 4, where higher-resolution units I, J, M and N are averaged, and encoded and transmitted in a lower-resolution grid, thus allowing the receiver to recreate higher-resolution unit J by means of the fractional overlap between the lower-resolution units of frames 0 and 4.

Once the pattern of fractional shifts applied over the sequence of frames 0 to 5 has been completed, the full higher-resolution grid can be reconstructed at the receiver.

It will be appreciated however that the above is only one possible example. In other implementations, different ratios of higher to lower resolution unit sizes may be used, and/or other shift patterns may be used. For example, another possible shift pattern requiring only a four-frame cycle would transmit:

Av (B,C,F,G)
Av (E,F,I,J)
Av (J,K,N,O)
Av (G,H,K,L)

In one embodiment, it is not necessary for there to be actual movement of the camera or object. Instead, the encoder may generate an indication of an artificial shift or pattern of shifts to be applied at the decoder to recreate the higher-resolution. That is to say, the “movement” may only be artificially generated for the sole purpose of reducing transmitted bitrate.

Alternatively, the shift may be based on actual movement. In this case, the movement may be detected by using gyro sensors of a mobile terminal in which the camera is housed so as to detect movement of the camera, or by using motion estimation techniques to detect movement of the scene or object being captured.

Along with the encoded lower-resolution units of frames 0 to 5, the encoder also transmits some side information in the encoded bitstream indicative of the shift or pattern of shifts to be applied as part of this scheme. This indication could take the form of a separate shift indication for each frame in the sequence; or more preferably in the case of an artificially generated shift the indication may take the form of a single indicator for the whole sequence, referring to a predetermined pattern of shifts to use for that sequence. In the latter case, the predetermined patterns may be pre-stored at both the encoder at both the transmitter side and the receiver side. For example, the codec may be operable in one or more different super-resolution modes defined by different respective shift patterns and/or resolution ratios (and preferably the codec will also have a more conventional mode not using super-resolution). The nature of the different modes will be understood by both the encoder at the transmitter side and the decoder at the receiver side, and the side information signalled from the transmitter may indicate the mode that has been used at the transmit side by the encoder.

The present invention thus uniquely uses a super-resolution scheme to deliberately down-grade the resolution being transmitted in a video stream in order to reduce the bit-rate, and then reconstruct the higher-resolution image again at the receiver. Of course it is not possible to get “free data”—but the idea is to trade bitrate for reconstruction time, since the scheme will require multiple frames in order to reconstruct the higher resolution image at the receiver, thus taking a longer time to obtain the higher resolution than if the data was simply transmitted at the higher resolution in each frame.

For this reason, the above-described feature may not be suited to very fast motion, though it may be useful for encoding motion which is slower but more detailed. In a particularly advantageous embodiment, because the blocks are encoded on a block-by-block basis, it is possible to encode different regions of the same video differently. E.g. based on motion estimation analysis of the video image, a slow moving background may be encoded using lower-resolution units, whilst a faster moving foreground in the same image may be encoded using a higher-resolution; or even vice versa. In this case the shift information may be signalled on a block-by-block basis, where each block comprises multiple lower-resolution units (or on a macroblock-by-macroblock basis, etc.).

It is particularly preferred to use this idea in conjunction with the global block list described above. That is, some

frames or some blocks or areas within a frame may be encoded using the global block list feature described above, whilst other frames or even other blocks or areas within the same frame may be encoded using the super-resolution feature described in this section. For example, the global block list could be used to encode blocks in areas that are relatively static whilst the super-resolution could be used to encode blocks in other areas where more detailed motion is occurring (so as to reduce a peak in bitrate that such motion might otherwise cause); or the global block list could be used to encode the faster motion of large objects whilst the super-resolution feature could be used to encode areas where less motion is occurring (because it is less suited to fast motion due to the time required to reconstruct the higher-resolution image).

In other embodiments the global block list could alternatively be used to signal reference blocks for encoding and decoding video using a more conventional super-resolution approach for increasing resolution beyond the intrinsic resolution of the camera.

In addition to the scaling and rotation, the above allows an implementation of a super resolution approach in compilation of the frequency domain and spatial domain algorithm. In this approach every other frame is moved by $\frac{1}{2}$ or $\frac{1}{4}$ of a pixel in a specific pattern that can be communicated to the decoder. While encoding in such a way, it is possible to derive a benefit of following minor motion shifts of $\frac{1}{2}$ or $\frac{1}{4}$ of a pixel or unit by simply finding matches in the direction of the motion which would otherwise have been missed. Additionally the reconstruction can be done in the spatial domain via pixel re-sampling. FIG. 14a shows an example in which the $\frac{1}{2}$ pixel shift between the two frames allows for four new pixels (A, B, C and D) to be created out of one pixel in the original frame.

The shift direction can come from the acquisition system as an encoding input or created to reduce the bitrate as a reverse scalability, e.g. sending CIF resolution instead of VGA. Perceptual Model

The present invention provides a technique for performing one or more aspects of video coding such as quantization, intra prediction coding or inter prediction coding in dependence on a perceptual model taking into account human sensitivity to data in the video signal.

As shown in the sketches of FIGS. 15a, 15b and 15c, a human being experiences a varying sensitivity in vision at different spatial and temporal frequencies. That is to say, the human eye can detect a finer contrast to light intensity (luminance) and colour (e.g. chrominance U and V) at mid-range spatial and temporal frequencies, whilst can distinguish only much coarser differences in luminance or colour at higher spatial and temporal frequencies (and to some extent may also experience less sensitivity at lower frequencies).

FIG. 15a shows a three-dimensional curve (or surface) taking into account both spatial frequency and temporal frequency. FIG. 15b shows a cross section through FIG. 15a at some given temporal frequency, and FIG. 15c shows a cross section through FIG. 15a at some given spatial frequency.

Luminance contrast can be defined as $\Delta L/L$, where ΔL is some change in luminance at some luminance L . Colour contrast can be defined as $\Delta E/E$, where ΔE is some change in colour at some colour value E . Luminance and colour contrast sensitivity are measures of the minimum ΔL and ΔE that can be detected by a person at L and E respectively—i.e. so a smaller minimum ΔL implies a higher contrast sensitivity, and similarly for ΔE .

An example of a contrast sensitivity curve and metrics for use in measuring human sensitivity can be found in: “DCT-

based Video Quality Evaluation”, Feng Xiao, Winter 2000, http://compression.ru/video/quality_measure/vqm.pdf.

Further examples of metrics for measuring human perceptual sensitivity are discussed in: “Design Considerations for Computationally Constrained Two-Way Real-Time Video Communication”, Bivolarski et al, Applications of Digital Image Processing XXXII, edited by Andrew G. Tesher, Proc. Of SPIE Vol. 7443, 2009.

The actual contrast sensitivity response curve for use in any given application can be found by empirically by surveying a large sample of human individuals, determining from them the minimum change in luminance and colour which each can detect in response to being shown various test images. The sensitivity response curve will be subjective for any given individual, and the results may vary for different sample groups.

According to a further particularly preferred aspect, the present invention uses information on human sensitivity to create non-regular quantization levels in the quantizer 4. That is, the invention may use different size quantization bins according to a quantization bin size distribution determined from empirical information on human perceptual sensitivity, such that video data to which the human viewer is more sensitive is quantized using smaller bins (more accurate, so less distortion); whilst video data to which the human viewer is less sensitive is quantized using larger bins (less accurate, so more distortion). For example, in the transformed frequency domain representation of each block, the quantizer 4 may quantize the frequency domain coefficients of the block which correspond to more sensitive frequencies with finer bin sizes, and quantize the frequency domain coefficients of the block which correspond to less sensitive frequencies with coarser bin sizes. The advantage is that the “bit budget” of the encoded bitstream can be spent encoding those coefficients to which a human is most perceptually sensitive, whilst avoiding spending bits on coefficients to which a human is less sensitive. The variation in bin size with spatial and/or temporal frequency is related to an empirically-determined curve of human perceptual sensitivity as sketched schematically in FIGS. 15a-15c.

In further embodiments, the quantizer 4 may be configured with non-regular quantization bin sizes in dependence on the colour or the magnitude of the luminance being quantized. E.g. if the human eye is more sensitive to mid range luminance and not so much to higher or lower luminance, then the mid range may be quantized with finer quantization levels than the extremes of luminance. Similarly, if the human eye is more sensitive to certain colours than others, then the quantizer 4 may quantize those colours with finer quantisation levels than the others.

In one particular example, as shown in FIG. 15d, the human eye may have different sensitivity to different colour ranges, and hence the quantizer may also use different quantization bin size distributions for different colour channels (and/or preferably different distributions for the luminance and chrominance channels).

Referring to FIG. 15e, this shows a representation of the possible DCT terms in a 8x8 frequency domain representation, with horizontal wavenumber k_x and vertical wavenumber k_y . Any given block may be represented as a sum of one or more of these 64 possible DCT terms. However, the human eye is less sensitive to those terms in the bottom right with relatively high values of k_x and/or k_y , marked (*) in FIG. 15e.

Note that as shown in FIG. 15g, this behaviour of the quantizer 4 is made possible because in preferred embodiments the input data to the quantizer 4 is the unpredicted frequency domain coefficients—i.e. the quantizer is placed

before the intra and inter prediction encoding in the encoder chain. Hence the quantizer has access to the actual absolute coefficients representing what will be seen by the recipient, and thus the quantizer 4 is able to adapt in relation to an estimated perception of the recipient. This is unlike the standard arrangement as shown in FIG. 7, in which the input to the quantizer 4 is the transformed residual data in which the perceptually relevant information has already been lost. Since the absolute values cannot be known from the residual data, the quantizer 76 can have no knowledge of the actual colours or luminance of the data being quantized.

An alternative arrangement would be that shown in FIG. 15h, whereby the quantizer 4 is placed after the prediction 6, 8 but with information on the absolute values being passed from the output of the transform stage 2 to the quantizer 4 via a separate control input.

In further embodiments, the quantization bin size distribution of the quantizer 4 may be made dependent on a structural metric of the block or blocks being encoded. An example of such a metric would be the Structural Similarity Index Metric (SSIM).

According yet a further possible aspect of the present invention, a perceptual model can alternatively or additionally be employed in the block matching process during intra and/or inter prediction coding. In this case, the prediction coding module 6, 8 is configured with information on human perceptual sensitivity, which it can use to modify the target block and/or reference block so as to reduce the bitrate incurred by the residual data, but at the same time with the modification being within a limit that will not be significantly perceived by the viewer. For example, the prediction module 6, 8 may modify the luminance and/or chrominance coefficients of the target block so as to match it more closely to the reference block as long as the modification is within a certain degree dependent on the perceptual model. This is unlike conventional prediction which is essentially colour-blind.

There are two possible modifications which reduce the bitrate in the final encoded stream. One is to reduce the magnitude of the difference between the coefficients of the target block and the coefficients of the reference block, because a smaller residual will require fewer bits to encode. Another is to reduce the variation in difference between the coefficients, because a residual with less variation will have less entropy and therefore require fewer bits when encoded by the entropy encoder.

One particular example is shown schematically in FIG. 15h. This shows the absolute magnitude of a few luminance coefficients at different spatial frequencies k_x . If the prediction module 6, 8 has information on the maximum perceptually noticeable or tolerable modification that can be applied to each coefficient—shown by error bars in FIG. 15h—then it may select to modify the luminance values to be more similar or in this particular example all four coefficients could be modified to be the same. A similar process can be applied to the chrominance channels. The maximum tolerable modification is dependent on spatial and/or temporal frequency as shown in FIGS. 15a-15d, and hence the maximum modification is different for different coefficients.

In further embodiments, the modification applied by the prediction module 6, 8 may be dependent on the colour or the magnitude of the luminance of the coefficients being encoded. E.g. the human eye may be more sensitive to mid range luminance and not so much to higher or lower luminance, or is more sensitive to certain colours than others.

In further embodiments, the modification may be made dependent on a structural metric of the block or blocks being encoded, e.g. the Structural Similarity Index Metric (SSIM).

In particularly preferred embodiments, the operation of the quantizer 4 and/or prediction coding module 6,8 is made dependent on a parameter which affects human perception.

One example of such a parameter is motion occurring in the video, either because the camera has moved or an object being filmed has moved. This motion can be detected by a motion estimation algorithm or by sensors in a mobile terminal in which the camera is mounted, e.g. gyro sensors or accelerometers. Humans are typically less sensitive to higher spatial frequencies in objects that are moving faster, so this information can be used to adapt the quantization bin size distribution of the quantizer 4, or to vary the degree of modification allowed to be applied by the prediction module 6,8.

Other non-limiting examples of parameters affecting human perception include the aspect ratio of the recipient user's screen, the size of the recipient user's screen, and the resolution of the recipient's screen. So this information can be used to adapt the quantization bin size distribution of the quantizer 4, or to vary the degree of modification allowed to be applied by the prediction module 6,8. To this end, the decoder at the recipient terminal may be arranged to feedback information on the parameter to the encoder at the transmitting terminal for use in dynamically adapting the quantization and/or prediction behaviour of the encoder (though in other embodiments the information could be inserted manually by the transmitting user based on his or her own knowledge of the recipient user's terminal). In one particular embodiment, the distance of the recipient user from his or her screen may be measured automatically by an auto-focus feature the recipient's webcam, and then fed back to the transmitter for using in adapting the quantization and/or prediction.

For any one or more of these parameters, preferably the parameter can be determined and updated dynamically during encoding of a given video signal, and may be fed back dynamically if appropriate. This means the quantizer 4 and/or prediction module 6,8 will switch "on the fly" between different perceptual models from one frame to the next, or even for different blocks within the same frame.

For instance, different perceptual sensitivity curves or bin size distributions may be stored for different amounts of on-screen motion or different distances of the recipient from the screen.

In preferred embodiments, the perceptual algorithm uses block matching prediction based on a perceptual model to establish the perceptual importance of a block and block matching contribution to the final rate. The block selection is performed within in a sorted array based on importance indexing. Selection of matched is done based on an aggregate from the indexes. The group of candidate blocks in the selections is then evaluated for perceptual quality and rate contribution. The best candidate is selected out of the lowest rate contributor that is matching or the most exceeding a quality goal block. Rate adjustment is done based on a degree of freedom estimated from the perceptual model in the frequency domain.

As discussed above, the use of a perceptual model gives the possibility to change the coefficients of each block, which can be altered to reduce the bit rate while keeping the perceptual quality intact. Another benefit of the algorithm is that it can be applied as an addition to current coding schemes as an additional avenue of reducing the bit rate as a lossless component of the algorithm in the frequency domain, however in this case the performance will be reduced.

The block similarity is determined from both the point of view of reduction of bit rates and human perception.

Both similarities can be used in combination. Such similarities may explore redundancies and luminance, colour, spatial, temporal and structural properties of video perception and sensing.

In the case where the transform is the first step of the codec algorithm, the perceptual based quantizer can be applied several steps before the temporal prediction and after. In the case of using a multidimensional transform where one of the components is a temporal part, the temporal quantization can be applied right after the transform.

There are preferably three basic steps in the perceptual model—luminance and chrominance interdependent sensitivity that is applied simultaneously on all color channels; spatial and temporal sensitivity that are interdependent and applied at the same time; and structural sensitivity that reflects the interdependency of between the spatial frequencies.

Additional aspects of the perceptual coding method are the ability of human vision system to distinguish spatial frequencies based on its spatial angular limitation or the distance from the physical display surface and the size and aspect ratio of the display. These factors will determine additional limitations on the ability of the display to produce the information that is being sent or on the human vision system to detect the some of the information that is sent in the encoded stream. Additional presentation conditions will be adjusted by receiving additional information from the stream encoder on both sides of the video conferencing peers, like distance from the screen of the viewer on the far end as well as light conditions like color temperature and scene overall brightness.

The block prediction in the frequency domain will preserve the information necessary to for the perceptual method.

The method presented here is unique based on the exclusive block prediction in a frequency domain that allows for proper integration of the aspects of the human vision system into the coding process.

Thus the sensitivity measure will aid the encoding process by removing all of the information from the stream that is not perceived from the human vision system.

Additionally, device dependant adjustment will provide additional rate savings where the device is not capable or displaying.

The motion dependant quantization will reduce the amount of bits based on the temporal contrast sensitivity while the objects are moving as well as taking into account the frame rate contribution. The human vision system has reduced contrast sensitivity for moving objects.

The luminance and chrominance interdependent sensitivity is nonlinear and is capable of detecting significantly lower number of colors the currently used 16.7 million by encoding in 24 bit. By implementing nonlinear color representation we will benefit of representing lower number of colors.

The distance between quantisation indices may be set according to the multidimensional quantisation model and generally would follow the Weber-Fechner Law. Model parameters may be used to look up quantisation values from a quantisation look-up table or tree.

The prediction is done within a sorted list of blocks. The search area within the list of predictor blocks will increase or decrease based on the sensitivity measure related to the perceptual sensitivity model and the viewing characteristics listed above.

The coefficients of the matched block may be modified to reduce the residual based on the perceptual model.

The following parameters may be used as variables in the perceptual model

the colour of the components of each block being encoded
the luminance of the components of the block being encoded
the colour of the components of the surrounding blocks
the luminance of the components of the surrounding blocks
the speed that a video object is moving
distance of user from screen
background light in the environment of the viewer
display characteristics like color, resolution and aspect ratio.

These models may be used alone or in combination or part of a multidimensional perceptual model.

The degree of compression may be adjusted in the quantization step and during matching and sorting. The complexity of the algorithms is reduced significantly by using sorting since it is a one pass per frame or slice and creates localized search based on block characteristics.

Referring again to FIG. 2, in embodiments the method may use block matching based on perceptual and objective metrics in the frequency domain to determine which blocks can be predicted. According to S2, each block $b_0 \dots b_{23}$ in the frequency domain may be assigned an index value based on objective and/or perceptual set of metrics. In Step S5, the candidate blocks may then be evaluated for closeness (a match) based on the metrics a perceptual threshold from the aggregate measure used to exclude definite differences that could not produce a match. The rate of the residual blocks and or objects is estimated as the use of perceptual model in this case creates an opportunity for further reduction of bit rate without loss of perceptual quality.

Referring again to FIG. 4, the spatial part of the transform can be adapted for each colour plane and applied in combination with the rest of the transform. The forward quantization stage 2 applies quantization based on video presentation and sensing properties. The method allows for further manipulation of the data of each block based on the degree of freedom which the perceptual model allows. The first two stages take out of the stream information that is irrelevant to presentation or sensing which can be determined in a real time and interactive fashion. In this way, the irrelevant components of the signal are not transmitted further into the chain of processing, thus reducing the complexity.

This method can further reduce the rate beyond quantization using the perceptual model. The quantized coefficients can be altered based on the perceptual tolerance interval for each specific point in the frame in which the perception remains constant and the block values can be varied to reduce the bit rate. This way in comparison with other encoding algorithms the proposed method can achieve better bit rate without affecting the human perception.

Implementation

The encoder elements 2, 4, 6, 8 and 10; and the decoder elements 12, 14, 16, 18 and 20 are each preferably implemented in software modules stored on a storage medium such as a hard drive or flash memory and arranged for execution on a processor. However, in other embodiments some or all of these elements could be at least partially implemented in dedicated hardwired circuitry.

In one possible embodiment, the invention may be implemented as an add-on to an existing encoder such as ISO standard H.264. That is, the input to the quantizer 4 in FIG. 4 will be an output from a standard encoder such as an H.264 encoder.

It will be appreciated that the above embodiments have been described only by way of example.

For instance, note that whilst the term "block" is used herein, in its most general sense this is not intended to imply and particular size, shape or level subdivision. It will be appreciated that in different schemes then various different divisions and subdivisions may be referred to by terms such as macroblock, block and sub-block, etc., but that the term "block" as used most generally herein may correspond to any of these or indeed any other constituent image portion being a division of a video frame.

Further, whilst the above has been described with reference to the example of a Discrete Cosine Transform into the frequency domain, it will be appreciated that other transforms such as the KLT or others can be used (some of which may not represent the transform domain in terms of spatial frequency coefficients but in terms of some other transform domain coefficients).

Further, whilst the above has been described in terms of a residual representing the subtracted difference between the coefficients of the target block and the coefficients of the reference block, this is not the only possibility for encoding the coefficients or values of the target block relative to those of the reference block. In other possible embodiments for example, the difference may be represented and signalled in terms of parameters of a correlation between the target block and the reference block such that the target can be predicted from the correlation, or in terms of coefficients of a filter that may be applied to the reference block to predict the target block. In these cases the prediction may not necessarily be lossless as in the case of a subtractive difference, but may instead be lossy such that the difference does not represent the exact difference. The term "difference" as used herein is not limited to subtractive difference nor to an exact difference.

Further, the present invention is not limited to implementation in any particular standard nor as an add-on to any particular standard, and may be implemented either as a new stand-alone codec, an add-on to an existing codec, or as a modification to an existing codec.

Other variants may be apparent to a person skilled in the art given the disclosure herein. The invention is not limited by the described embodiments, but only by the appendant claims.

The invention claimed is:

1. A method comprising:

receiving a video signal comprising a plurality of frames, each frame comprising one or more portions of video data;

transforming each of said portions into a transform domain, so as for each portion to generate transformed video data comprising a set of coefficients, each coefficient representing a contribution to a property of the portion at a different spatial or temporal frequency;

quantizing the video data of each portion according to a quantization bin size distribution, the quantization bin size distribution varying dynamically during ongoing encoding such that a different quantization bin size distribution is applied to different transformed video data, the quantization bin size distribution being dependent on: a structural metric of the one or more portions of the video data and a measure of human sensitivity to the video data such that a portion of the video data to which a human viewer is more sensitive is quantized using a smaller quantization bin size and a portion of the video data to which the human viewer is less sensitive is quantized using a larger quantization bin size, the quantization bin size distribution varying dynamically according to information received automatically from a recipient terminal; and

41

encoding the video data of each portion based, at least in part, on the quantizing.

2. The method of claim 1, wherein:

said quantization comprises quantizing the coefficients of each portion; and

the quantization bin size distribution uses different quantization bin sizes for different colour properties of the transformed video data.

3. The method of claim 1, wherein the quantization bin size distribution is related to a measure of human sensitivity to said property at different spatial frequencies.

4. The method of claim 1, wherein the quantization bin size distribution uses different quantization bin sizes for different ones of said coefficients.

5. The method of claim 1, wherein the property comprises one of a luminance and a colour property.

6. The method of claim 5, wherein some of said coefficients represent luminance and some of said coefficients represent at least one colour property, and the method further comprising using a different quantization bin size distribution for quantizing the luminance and for quantizing the at least one colour property.

7. The method of claim 5, wherein some of said coefficients represent a first colour channel and some of said coefficients represent a second colour channel, and the method further comprising using different quantization bin size distributions for the first and second colour channels.

8. The method of claim 1, further comprising adapting the quantization bin size distribution according to information identifying spatial frequencies in objects included in the video data.

9. The method of claim 1, wherein the quantization bin size distribution is variable, during encoding, in dependence on a distance of the human viewer from a screen of the recipient terminal.

10. The method of claim 1, wherein the quantization bin size distribution is dynamically varied during ongoing encoding in dependence on a parameter which affects human perception.

11. The method of claim 10, wherein the bin size distribution is varied dynamically in dependence on motion in the video signal, the motion being a parameter which affects human perception.

12. The method of claim 11, wherein the motion is determined by a motion estimation algorithm applied to the frames of the video signal.

13. The method of claim 11, wherein the motion is detected by physical sensors in a mobile terminal housing a camera which generates said video signal.

14. The method of claim 9, comprising transmitting an indication of the bin size distribution to a decoder for use in decoding the video signal.

15. The method of claim 1, wherein the video data is input to the quantizer before inter or intra encoding such that the video data represents absolute values.

16. A method comprising:

receiving a video signal comprising a plurality of frames, each frame comprising one or more portions of video data;

transforming each of said portions into a transform domain, so as for each portion to generate transformed video data comprising a set of coefficients, each coefficient representing a contribution to a property of the portion at a different spatial or temporal frequency;

quantizing the video data of each portion, the quantizing including quantizing the coefficients of each portion;

42

encoding the video data of each portion according to a quantization bin size distribution, the quantization bin size distribution varying dynamically during ongoing encoding such that a different quantization bin size distribution is applied to different transformed video data, the quantization bin size distribution being dependent on: a measure of human sensitivity to the video data such that a portion of the video data to which a human viewer is more sensitive is quantized using a smaller quantization bin size and a portion of the video data to which the human viewer is less sensitive is quantized using a larger quantization bin size, the encoding including using different quantization bin size distributions for different colour channels; and

transmitting the encoded bitstream to a decoder, the encoded bitstream including side information indicating a prediction method to be used during decoding, the prediction method being dependent on the measure of human sensitivity to the video data.

17. An encoder comprising:

a signal processing apparatus configured to implement a method of encoding a video signal for transmission to a decoder of a receiving terminal, the method comprising: receiving a video signal comprising a plurality of frames, each frame comprising one or more portions of video data;

transforming each of said portions into a transform domain, so as for each portion to generate transformed video data comprising a set of coefficients, each coefficient representing a contribution to a property of the portion at a different spatial or temporal frequency;

quantizing the video data of each portion according to a quantization bin size distribution, the quantizing including quantizing the coefficients of each portion, the quantization bin size distribution varying dynamically during ongoing encoding such that a different quantization bin size distribution is applied to different transformed video data, the quantization bin size distribution being dependent on: a structural metric of the one or more portions of the video data and a measure of human sensitivity to the video data such that a portion of the video data to which a human viewer is more sensitive is quantized using a smaller quantization bin size and a portion of the video data to which the human viewer is less sensitive is quantized using a larger quantization bin size, the quantization bin size distribution varying dynamically during ongoing encoding in dependence on information received automatically from a recipient terminal;

encoding the video data of each portion based, at least in part, on the quantizing.

18. The encoder of claim 17, wherein the quantization bin size distribution is related to a measure of human sensitivity to said property at different spatial or temporal frequencies.

19. The encoder of claim 17, wherein the quantization bin size distribution uses different quantization bin sizes for different ones of said coefficients.

20. The encoder of claim 17, wherein the property comprises one of a luminance and a colour property.

21. The encoder of claim 20, wherein some of said coefficients represent luminance and some of said coefficients represent at least one colour property, and the method further comprising using a different quantization bin size distribution for quantizing the luminance and for quantizing the at least one colour property.

43

22. The encoder of claim 20, wherein some of said coefficients represent a first colour channel and some of said coefficients represent a second colour channel, and the method further comprising using different quantization bin size distributions for the first and second colour channels.

23. The encoder of claim 17, wherein the quantization bin size distribution is further dependent on an aspect ratio, a size, and a resolution of a screen that is to present the video data after the quantizing.

24. The encoder of claim 17, wherein the quantization bin size distribution is variable in dependence on a parameter which affects human perception and a distance of the human viewer from a screen of the receiving terminal being a parameter which affects human perception.

25. The encoder of claim 17, wherein the quantization bin size distribution is dynamically varied during ongoing encoding in dependence on a parameter which affects human perception.

26. The encoder of claim 25, wherein the bin size distribution is varied dynamically in dependence on motion in the video signal, the motion being a parameter which affects human perception.

27. The encoder of claim 26, wherein the motion is determined by a motion estimation algorithm applied to the frames of the video signal.

28. The encoder of claim 26, wherein the motion is detected by physical sensors in a mobile terminal housing a camera which generates said video signal.

29. The encoder of claim 24, wherein the method further comprises transmitting an indication of the bin size distribution to the decoder for use in decoding the video signal.

30. The encoder of claim 17, wherein the video data is input to the quantizer before inter or intra encoding such that the video data represents absolute values.

31. An encoder comprising:

a signal processing apparatus configured to implement a method of encoding a video signal for transmission to a decoder of a receiving terminal, the method comprising: receiving a video signal comprising a plurality of frames, each frame comprising one or more portions of video data;

transforming each of said portions into a transform domain, so as for each portion to generate transformed video data comprising a set of coefficients, each coefficient representing a contribution to a property of the portion at a different spatial or temporal frequency, quantizing the video data of each portion, the quantizing including quantizing the coefficients of each portion; and

encoding the video data of each portion according to a quantization bin size distribution, the quantization bin size distribution varying dynamically during ongoing encoding such that a different quantization bin size distribution is applied to different transformed video data, the quantization bin size distribution being dependent on: a structural metric of the one or more portions of the video data and a measure of human sensitivity to the video data such that a portion of the video data to which a human viewer is more sensitive is quantized using a smaller quantization bin size and a portion of the video

44

data to which the human viewer is less sensitive is quantized using a larger quantization bin size, the quantization bin size distribution varying dynamically according to information received automatically from a recipient terminal.

32. The encoder of claim 31, wherein the quantization bin size distribution is related to a measure of human sensitivity to said property at different spatial or temporal frequencies.

33. The encoder of claim 31, wherein the quantization bin size distribution uses different quantization bin sizes for different ones of said coefficients.

34. The encoder of claim 31, wherein the property comprises one of a luminance and a colour property.

35. The encoder of claim 34, wherein some of said coefficients represent luminance and some of said coefficients represent at least one colour property, and the method further comprising using a different quantization bin size distribution for quantizing the luminance and for quantizing the at least one colour property.

36. The encoder of claim 34, wherein some of said coefficients represent a first colour channel and some of said coefficients represent a second colour channel, and the method further comprising using different quantization bin size distributions for the first and second colour channels.

37. The encoder of claim 31, wherein the quantization bin size distribution is further dependent on an aspect ratio, a size, and a resolution of a screen that is to present the video data after the quantizing.

38. The encoder of claim 31, wherein the quantization bin size distribution is variable in dependence on a parameter which affects human perception and a distance of the human viewer from a screen of the recipient terminal being a parameter which affects human perception.

39. The encoder of claim 31, wherein the quantization bin size distribution is dynamically varied during ongoing encoding in dependence on a parameter which affects human perception.

40. The encoder of claim 39, wherein the bin size distribution is varied dynamically in dependence on motion in the video signal, the motion being a parameter which affects human perception.

41. The encoder of claim 40, wherein the motion is determined by a motion estimation algorithm applied to the frames of the video signal.

42. The encoder of claim 40, wherein the motion is detected by physical sensors in a mobile terminal housing a camera which generates said video signal.

43. The encoder of claim 38, wherein the method further comprises transmitting an indication of the bin size distribution to the decoder for use in decoding the video signal.

44. The encoder of claim 31, wherein the video data is input to the quantizer before inter or intra encoding such that the video data represents absolute values.

45. The method of claim 16, wherein the property of the portion includes a luminance property or a colour property.

46. The method of claim 16, wherein the quantization bin size distribution is variable in dependence on a distance of the human viewer from a screen of the recipient terminal.

* * * * *